arXiv:1409.1551v1 [cs.IT] 4 Sep 2014

Synchronizing Edits in Distributed Storage Networks

Salim El Rouayheb, Sreechakra Goparaju, Han Mao Kiah, Olgica Milenkovic

September 5, 2014

Abstract

We consider the problem of synchronizing data in distributed storage networks under an edit model that includes deletions and insertions. We present two modifications of MDS, regenerating and locally repairable codes that allow updates in the parity-check values to be performed with one round of communication at low bit rates and using small storage overhead. Our main contributions are novel protocols for synchronizing both hot and semi-static data and protocols for data deduplication applications, based on intermediary permutation, Vandermonde and Cauchy matrix coding.

1 Introduction

Coding for distributed storage systems has garnered significant attention in the past few years [1–4], due to the rapid development of information technologies and the emergence of Big Data formats that need to be saved or disseminated in a distributed fashion across a network. Reliable and efficient storage, transfer, retrieval and update of distributed data are challenging computational and networking tasks in which coding theory plays an indispensable role. One of the most significant applications of coding methods is to ensure data integrity under distributed component and node failure, via controlled redundancy introduced at different locations. This is accomplished by implementing modified Reed-Solomon, regenerating and local repair coding solutions [5–10] that extend and generalize the traditional coding paradigms that are part of CDs, DVDs, flash memories and RAID like storage systems [11–13].

Typical distributed storage systems (DSS) such as Google's Bigtable, Microsoft Azure and the Hadoop Distributed File System, are all designed to scale to very large file sizes, to allow for low-latency data access and to store redundant coded copies of files over a set of servers, disks or nodes connected via a communication network. Two key functionalities of codes for DSS systems are reconstructability of files via access to a subset of the nodes; and content repairability of failed nodes that does not compromise data reconstruction capabilities. Evidently, both these functionalities need to be retained when the content of files undergoes *edits* which arise in dynamically changing storage systems and networked systems with load asymmetries involving "hot data" [14] or data deduplication features [15]. Other examples in which frequent edits are encountered are Dropbox and Sugarsync shared file systems [16], in which large number of user files are being independently updated at multiple locations. Current solutions for synchronization protocols either assume that data in uncoded or do not fully exploit the distributed nature of information, while deduplication methods mostly apply to read-only architectures and are in an early stage of development as far as distributed systems are concerned [17, 18].

When files are edited or deduplicated, the changes in the information content need to be communicated to the redundant storage nodes so that the DSS retains its reconstruction and repair capabilities, and so that the communication and update costs are minimized. This is a challenging task, as edits such as deletions and insertions, which most commonly arise in practice, cause synchronization issues that appear hard to account for in the DSS encoded domain. In Dropbox and related systems, deletion and insertion synchronization issues are resolved in the *uncoded* domain via the use of the well known rsync [19] and dsync [20] or zsync (http://zsync.moria.org.uk/) algorithms, related to a number of file synchronization methods put forward in the information theory literature [21–28]. There, uncoded stored copies of a file are synchronized from an edited user's copy. More specifically, the problems involve a single user and a single node storing a replicate of the user's file. After the user edits his/her file, assuming *no knowledge* of the edits, the user and the storage node communicate interactively until their files are matched or until one node matches the master copy of the user. In the DSS scenario we propose to analyze, one may assume both *full knowledge* of the edits or unknown edits, pertaining to the scenarios that edits were made by the users themselves or by another user sharing the file, respectively. The core problem in this case is to efficiently update *coded* copies in storage nodes with minimal communication rate. It is worth pointing out that this type of question regarding synchronization from edits in DSS systems fundamentally differs from update efficient codes studied in [29–32]. In these contributions, the authors minimize the *number of nodes* that need to be updated when a user's file is changed. Consequently, the edits in consideration may be vaguely viewed as substitutions, in which case, minimizing the communication cost reduces to minimizing the number of storage nodes a user needs to recruit during update. In contrast, this line of work is concerned with edits of the form of insertions or deletions, where such simplified reductions do not apply. And although one may view a deletion as a sequence of substitution errors, using update efficient codes with such preprocessing is highly suboptimal; in the worst case, one deletion is equivalent to ℓ substitutions, where ℓ is the file length. Furthermore, instead of minimizing the number of storage nodes a user needs to communicate, the objective of synchronizing from deletions and insertions is to minimize the communication cost between the users and storage node even at the cost of introducing a small, controlled amount of storage overhead.

The contributions of the paper are three-fold. First, we describe a number of edit models and accompanying coding protocols for efficient synchronization in a distributed storage environment that maintains regenerative properties. The synchronization protocols are based on a simple new scheme termed *intermediary coding*, which flexibly changes the structure of the code so as to allow reduced communication complexity between nodes in a distributed system. The intermediary coding schemes also offer flexibility in terms of *accommodating a very broad family of coding schemes* used in storage devices (such as erasure codes, regenerating codes, locally repairable codes etc). Second, we study extensions of the edit models related to deduplication applications and different data types, such as hot and semi-static data. Third, we provide worst and average case communication cost analyses for different edit models and synchronization protocols. This analysis reveals that traditional schemes require a significantly higher communication cost than schemes based on intermediary coding, both in the worst case and average case scenario. This may be attributed to the fact that traditional encoding requires each node to communicate symbols in the *span* of all deletion positions in different nodes, while intermediary coding allows for reducing the communication cost to the number of bits needed to encode the particular edit only.

The paper is organized as follows. Section 2 introduces the relevant notation and provides

the precise problem statement. Section 3 contains the description of simple update protocols and their underlying communication cost when traditional DSS encoding methods are used. Section 4 contains our main result, a collection of encoding algorithm and protocols for data synchronization that have overhead communication cost that is a constant factor away from the fundamental limits derived in Section 5. Section 5 then examines both storage overhead and communication overhead and explains how to trade between these two system parameters. Section 6 provides a short discussion on how to handle unknown edit positions. The average case analysis of the communication cost of traditional and intermediary coding schemes, the description of hybrid and deduplication schemes are presented in the Appendix section of the paper.

2 Notation and Problem Statement

Throughout the paper, we use $[n]$ to denote the set of integers $\{1, 2, \ldots, n\}$. For a finite set X and $k \leq |X|$, the collection of k-subsets of X is denoted by $\binom{X}{k}$. The symbol $\mathbb{F}_q$ is reserved for a finite field of order q. The vector space of all vectors of length ℓ, the vector space of all matrices of dimensions $m \times n$, and the vector space of all tensors of dimensions $m \times n \times \ell$ are denoted by $\mathbb{F}_q^\ell$, $\mathbb{F}_q^{m\times n}$ and $\mathbb{F}_q^{m\times n\times \ell}$, respectively. For $i \in [\ell]$, x_i represents the ith coordinate of a vector $\mathbf{x} \in \mathbb{F}_q^\ell$, while $\mathbf{e}_i$ stands for the ith standard basis vector.

Given a matrix $\mathbf{M} \in \mathbb{F}_q^{m\times n}$ and a subset of the rows $R \subseteq [m]$ and a subset of the columns $C \subseteq [n]$, $\mathbf{M}|_{R\times C}$ represents the $|R| \times |C|$ matrix obtained by restricting $\mathbf{M}$ to rows in R and columns in C. We use analogous definitions for tensors.

We provide next a straightforward example that explains the underlying motivation for the work as well as the difficulties encountered in synchronizing coded data.

Example 1. Consider two users with data vectors $\mathbf{u}^{(1)} = \left(u_1^{(1)}, u_2^{(1)}, u_3^{(1)}, u_4^{(1)}, u_5^{(1)}\right)$ and $\mathbf{u}^{(2)} = \left(u_1^{(2)}, u_2^{(2)}, u_3^{(2)}, u_4^{(2)}, u_5^{(2)}\right)$, both of length five, over $\mathbb{F}_q$. Here, q is chosen based on the smallest number of consecutive, editable bits. Suppose we have three storage nodes where nodes 1, 2 and 3 store user information $\mathbf{u}^{(1)}$ and $\mathbf{u}^{(2)}$, and parity information $\mathbf{u}^{(1)} + \mathbf{u}^{(2)}$, respectively. Then the system is able to reconstruct both data blocks $\mathbf{u}^{(1)}$ and $\mathbf{u}^{(2)}$ and repair any failed node by accessing any two nodes.

Suppose both data blocks are subjected to a single symbol deletion, and that the resulting data blocks are $\tilde{\mathbf{u}}^{(1)} = \left(u_1^{(1)}, u_2^{(1)}, u_3^{(1)}, u_4^{(1)}\right)$ and $\tilde{\mathbf{u}}^{(2)} = \left(u_2^{(2)}, u_3^{(2)}, u_4^{(2)}, u_5^{(2)}\right)$. What protocol should the users employ and what information do they have to communicate to the three storage nodes so as to retain both reconstruction and repair functionalities, all with minimal data transmission cost?

One option is for the nodes to update their respective contents to $\tilde{\mathbf{u}}^{(1)}$, $\tilde{\mathbf{u}}^{(2)}$ and

$$\tilde{\mathbf{u}}^{(1)} + \tilde{\mathbf{u}}^{(2)} = \left(u_1^{(1)} + u_2^{(2)}, u_2^{(1)} + u_3^{(2)}, u_3^{(1)} + u_4^{(2)}, u_4^{(1)} + u_5^{(2)}\right).$$

Clearly, it is both necessary and sufficient for the user with data block $\tilde{\mathbf{u}}^{(1)}$ to communicate his/her deletion *position* to node 1. A similar statement holds true for the user with data block $\tilde{\mathbf{u}}^{(2)}$ and node 2. But what is the minimum communication complexity needed for node 3 to update its content to $\tilde{\mathbf{u}}^{(1)} + \tilde{\mathbf{u}}^{(2)}$? At first glance, it appears that both nodes 1 and 2 should transmit their whole content to node 3, as deletions occurred in their first and last position, respectively. The goal of this work is to show that, akin to the notion of functional repair in DSS [6], one may significantly

save in communication complexity by updating the content of node 3 through a flexible change in the code structure.

2.1 Coding for Distributed Storage Systems

Let $\mathbf{x} = \left(x^{(1)}, x^{(2)}, \ldots, x^{(B)}\right) \in \mathbb{F}_q^B$ be an information vector to be stored in a distributed storage system (DSS). We call each $x^{(s)}$, for $s = 1, \ldots, B$, a *data unit*. A DSS is equipped with an *encoding* function, as well as a set of reconstruction and a set of repair algorithms specifying a coding scheme. The encoding algorithm converts the vector $\mathbf{x}$ into n vectors of length α and stores them in n *storage nodes*, while the reconstruction algorithm recovers $\mathbf{x}$ from the contents of any $k \leq n$ out of n nodes. In addition, when a node fails, one uses the content of a subset of $d \leq n$ nodes to repair the content of the failed node.

More formally, we have the following definitions.

1. An *encoding* function is a map $\textsc{Encode} : \mathbb{F}_q^B \to \mathbb{F}_q^{n\times\alpha}$, where α is a given code parameter that has to be chosen so that the encoding function of interest is constructable.

2. For any $T \in \binom{[n]}{k}$, a map $\textsc{Reconstruct}(T) : \mathbb{F}_q^{k\times\alpha} \to \mathbb{F}_q^B$ is termed a *reconstruction* function if for all $\left(x^{(1)}, x^{(2)}, \ldots, x^{(B)}\right) \in \mathbb{F}_q^B$,

$$\textsc{Reconstruct}(T)\left(\textsc{Encode}\left(x^{(1)}, x^{(2)}, \ldots, x^{(B)}\right)\Big|_{T\times[\alpha]}\right) = \left(x^{(1)}, x^{(2)}, \ldots, x^{(B)}\right).$$

3. Given a $t \in [n]$ and $T \in \binom{[n]\setminus\{t\}}{d}$, a map $\textsc{Repair}(t, T) : \mathbb{F}_q^{d\times\alpha} \to \mathbb{F}_q^{\alpha}$ is termed an *exact repair* function if
$$\textsc{Repair}(t, T)(\mathbf{C}|_{T\times[\alpha]}) = \mathbf{C}|_{\{t\}\times[\alpha]},$$
where $\mathbf{C}$ is the information stored over n nodes.

Depending on the set of subsets T over which a repair function is defined on, one recovers the various descriptions of codes for DSSs studied in the literature:

1. *Maximum distance separable (MDS) codes* [33, Ch. 11]. Here, we have $d = k = B$ and $\alpha = 1$, and the reconstruction algorithm can also be used as the repair algorithm. We simply denote an MDS code of length n and dimension k by $[n, k]$, following the standard notation for the dimensions of a linear error-correcting code, where it is understood that the minimum Hamming distance of the code equals $n - k + 1$. One of the key properties of MDS codes exploited for repair is the fact that every k-subset of the n coordinates is an information set.

2. *Regenerating codes* [34, 35]. Here, $d \geq k$, and we require a repair function $\textsc{Repair}(t, T)$ for all $t \in [n]$ and all $T \in \binom{[n]\setminus\{t\}}{d}$. We note that in bandwidth-limited DSSs, it is of importance to define an intermediary function, mapping words from $\mathbb{F}_q^{d\times\alpha}$ to $\mathbb{F}_q^{\alpha}$. More precisely, a repair algorithm $\textsc{Repair}(t, T)$ maps inputs from $\mathbb{F}_q^{d\times\alpha}$ to $\mathbb{F}_q^{d\times\beta}$ and then to $\mathbb{F}_q^{\alpha}$, where β is usually smaller than α and indicates the required amount of downloaded information (bits). However, as we observe later, this intermediary function is not needed in our subsequent analysis.

3. *Locally repairable codes (LRC)* [9, 10, 36]. Here, $d < k$, and for all $t \in [n]$, we only require a repair function $\textsc{Repair}(t, T)$ for *some* $T \in \binom{[n]\setminus\{t\}}{d}$. Efficient repair is achieved by minimizing the number of nodes that needs to be contacted. In contrast to regenerating codes, we only require information from a particular set of d nodes to repair a node, and the amount of information downloaded is not considered.

We broadly refer to the three aforementioned families of encodings for DSSs as (n, k, d, α, B) DSS codes, and focus on code maps that are linear.

2.2 Problem Description

Assume next that the units of a data block are updated via deletion or insertion edits[1]. For $s \in [B]$, consider the set $N(s) \triangleq \{t \in [n] : \textsc{Encode}(\mathbf{e}_s)|_{\{t\}\times[\alpha]} \neq \mathbf{0}\}$; in words, $N(s)$ denotes the set of nodes that needs to be updated when the data unit indexed by s is edited. We say that the nodes in $N(s)$ are *connected to unit* s.

Given a (n, k, d, α, B) DSS code, one can extend it to a $(n, k, d, \alpha\ell, B\ell)$ DSS code for any choice of a positive integer ℓ. This extension is akin to data-striping in RAID systems, amounting to a simple conversion of a symbol to a vector of symbols. One can also view this construction as a means of dividing the $B\ell$ data units into ℓ groups of B symbols each and then applying to each group of symbols an (n, k, d, α, B) DSS encoding. Equivalently, in this $(n, k, d, \alpha\ell, B\ell)$ DSS coding scheme we may regard the given information as a B-tuple of *data blocks*, each of length ℓ, say $\left(\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \ldots, \mathbf{x}^{(B)}\right)$. The choice of ℓ is dictated by the storage system at hand, in particular, by the size of the block (which is a system parameter) and the expected size of an edit. Consequently, we assume that for algorithmic approaches that ℓ is a fixed system parameter. Nevertheless, for asymptotic analysis, we use the standard modeling assumption in which ℓ is allowed to grow arbitrarily large.

Example 1 (Continued). Recall the code of Example 1. There, we tacitly considered a single parity $[3, 2]$ MDS code, or a DSS code with parameters $n = 3$, $k = 2$, $d = 2$, $\alpha = 1$, $B = 2$ and $\ell = 5$, with codewords of the form $\left(u^{(1)}, u^{(2)}, u^{(1)} + u^{(2)}\right)$, and $u^{(1)}, u^{(2)}$ belonging to a finite field. Assume that there are two users with data blocks $\mathbf{u}^{(1)}$ and $\mathbf{u}^{(2)}$ of length ℓ. Then by having nodes 1, 2 and 3 store $\mathbf{u}^{(1)}, \mathbf{u}^{(2)}, \mathbf{u}^{(1)} + \mathbf{u}^{(2)}$, we obtain a $(3, 2, 2, \ell, 2\ell)$ DSS code. Furthermore, we observe that the connected nodes of the user with data block $\mathbf{u}^{(1)}$ are nodes 1 and 3, while the connected nodes of the user with data block $\mathbf{u}^{(2)}$ are nodes 2 and 3.

The edit model of interest assumes that the data blocks are subjected to deletions performed in an independent fashion by B different users[2]. More precisely, two different models for edits are studied, including:

1. *The uniform edits model,* in which each data block has the same number of deletions and thus the resulting data blocks all have the same length. Here the number of deletions is $o(\ell/\log \ell)$. This model is used to describe the main ideas behind the work in a succinct and notationally simple manner, but may not be of practical importance. The model is also amenable for combinatorial analysis.

[1]Consider a string $(x_1, x_2, \ldots, x_\ell)$. A *deletion* at position i results in $(x_1, x_2, \ldots, x_{i-1}, x_{i+1}, \ldots, x_\ell)$, while an *insertion* of symbol a at position i results in $(x_1, x_2, \ldots, x_{i-1}, a, x_i, \ldots, x_\ell)$.

[2]It may be possible that a user edits a number of different data blocks. However, for simplicity, our model assumes that each data block is edited by one user.

2. *The nonuniform edits model,* in which each data block has a possibly different number of edits. This model may be analyzed both in a combinatorial and probabilistic setting. In the former case, we assume $D \leq B\,\ell$ edits for all B users. In the latter case, we assume that one is given the probability p of deleting any particular symbol in a data block, resulting in an average number of $p\,\ell$ edits per data block. Note that p may depend on ℓ.

More generally, the problem of interest may be stated as follows: find the "best" protocol for the B users to communicate their edits to the storage nodes, so that the storage nodes can update their information while maintaining both reconstruction and repair functionalities. The word "best" may refer to the smallest communication cost, smallest required combination of communication and additional storage cost, etc. For simplicity, we first focus on the case of smallest communication cost, defined as the average number of bits transmitted from one user to all the storage nodes. An obvious suboptimal approach is for each user to send the entire data file of length ℓ to each connected storage node, so that each such node may update its information according to the encoding function at hand. However, as the number of edits made to a data block may be (and usually is) much smaller compared to the data block size, a communication cost of ℓ symbols or $\ell \log q$ bits may be highly suboptimal.

On the other hand, suppose a user performs a single edit. To encode the information about this edit, one requires $\log \ell$ bits for a deletion (for encoding the position) and $\log \ell + \log q$ bits for an insertion (to encode the position and the symbol). Hence, assuming that the data blocks and edits are uncorrelated, one requires each user to communicate at least $\log \ell + \log q$ bits for each edit. As a result, it is straightforward to see that a (loose) lower bound on the communication cost needed for synchronization from a *constant number* of deletions is $\log \ell$ (deletions) or $\log \ell + \log q$ (insertion) bits, as at least one user has to communicate the information about one of its edits. When the number of edits is of the order of the length of the data unit ℓ, this bound may not apply, as it may be more efficient to communicate "nonedits". This difficult issue will not be discussed in this work, although all schemes proposed in this work apply to this case as well, but without a proof of order-optimality. More detailed descriptions and tighter lower bounds on the communication complexity are given in Section 5.

In what follows, we propose a number of schemes that achieve a communication cost of $O(\log \ell + \log q)$ bits that is of the order of the intuitive lower bound. To facilitate such low communication cost, we introduce additional storage overhead needed to define an intermediary encoding function, which we refer to as intermediary coding. The gist of the encoding method is to transform the information, and hence the codes applied to data blocks, via permutation, Vandermonde and Cauchy matrices (the resulting schemes are subsequently referred to as Scheme P, V and C, respectively). The key property of the transforms is that they reduce the update and synchronization communication cost by changing the code structure. The storage overhead induced by this codes is carefully controlled by choosing the parameters of the corresponding matrices, as described in Section 5.2. In Section 5.2, we also demonstrate that our schemes are optimal in terms of storage allocation when the number of edits is $o(\ell/\log \ell)$. Under the same condition of $o(\ell/\log \ell)$ edits, we demonstrate in Appendix B that on the average our schemes outperform schemes that do not utilize intermediary encoding functions.

In general, our derivations and methods do not rely on specific assumptions on the network topology. For the proofs of lower bounds or fundamental limits, a user is allowed to communicate with any other user or storage node and vice versa, i.e., the storage network is a complete graph. However, users are naturally assumed to communicate only with a "minimal" set of storage nodes

	Assumptions	Schematic Diagram
Topology used for proving fundamental performance bounds.	A user is allowed to communicate with any other user or storage node and vice versa.	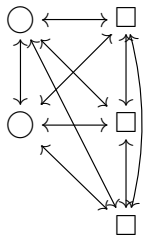
Traditional Coding Scheme (See Scheme T)	A user is allowed to communicate to any of his/her connected storage nodes and all users are allowed to have a two-way communication with a designated central node tasked with computing the span of all the deletions.	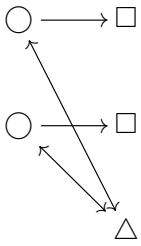
Intermediary Coding Schemes (See Scheme P and V)	A user is allowed to communicate to any of his/her neighboring storage nodes.	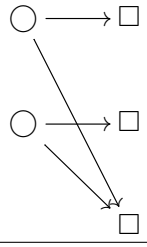

Here, ◯ denotes a user, while □ and △ denote a node and a designated central node, respectively.

Table 1: Communication Network Topology Models.

which contains encodings or systematic repetition of their data block symbols. These assumptions are summarized in Table 1.

3 Communication Cost with Traditional Encoding Schemes

As discussed earlier, assume that one is given an (n, k, d, α, B) DSS code over $\mathbb{F}_q$ with functions Encode, Reconstruct and Repair, and that for a given ℓ, this code is augmented to an $(n, k, d, \alpha\ell, B\ell)$ DSS code. Next, we state explicitly the encoding, reconstruction and repair functions for this scheme and examine the requisite communication between users and storage nodes under a single deletion model. The uniform deletion model may be analyzed in the same manner. We first show that in the worst case, the number of bits communicated is at least $(\ell - 1)$ symbols, or $(\ell - 1) \log q$ bits; we then proceed to introduce a scheme that lowers this communication cost to an order-optimal level.

More formally, for a fixed ℓ, in order to construct an $(n, k, d, \alpha\ell, B\ell)$ DSS code, we extend

ENCODE to the linear map[3] $\textsc{Encode}(\ell) : \mathbb{F}_q^{B\times\ell} \to \mathbb{F}_q^{n\times\alpha\times\ell}$, such that for $i \in [\ell]$,

$$\textsc{Encode}(\ell)\left(\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \ldots, \mathbf{x}^{(B)}\right)\Big|_{[n]\times[\alpha]\times\{i\}} = \textsc{Encode}\left(x_i^{(1)}, x_i^{(2)}, \ldots, x_i^{(B)}\right). \tag{1}$$

Hence, we regard the information stored at the n nodes as an $n \times \alpha \times \ell$ tensor. For $i \in [\ell]$, the ith slice of this tensor is the $n \times \alpha$ array $\textsc{Encode}\left(x_i^{(1)}, x_i^{(2)}, \ldots, x_i^{(B)}\right)$. We refer to this particular encoding $\textsc{Encode}(\ell)$ as a *traditional encoding* function. In the subsequent sections, we define another encoding function that uses $\textsc{Encode}(\ell)$ as a building block, and this novel encoding function will represent a crucial component of our low-communication cost synchronization schemes.

Next, we provide the corresponding reconstruction and repair functions accompanying traditional encoding and verify that we indeed have a DSS code.

In particular, we define $\textsc{Reconstruct}(T;\ell) : \mathbb{F}_q^{k\times\alpha\times\ell} \to \mathbb{F}_q^{B\times\ell}$ for $T \in \binom{[n]}{k}$ and repair function $\textsc{Repair}(t, T;\ell) : \mathbb{F}_q^{d\times\alpha\times\ell} \to \mathbb{F}_q^{\alpha\times\ell}$ for $t \in [n]$ and $T \in \binom{[n]\setminus\{t\}}{d}$ via

$$\textsc{Reconstruct}(T;\ell)\left(\mathbf{C}|_{T\times[\alpha]\times[\ell]}\right)\big|_{[B]\times\{i\}} \triangleq \textsc{Reconstruct}(T)\left(\mathbf{C}|_{T\times[\alpha]\times\{i\}}\right) \qquad \text{for } i \in [\ell], \tag{2}$$

$$\textsc{Repair}(t, T;\ell)\left(\mathbf{C}|_{T\times[\alpha]\times[\ell]}\right)\big|_{[\alpha]\times\{i\}} \triangleq \textsc{Repair}(t, T)\left(\mathbf{C}|_{T\times[\alpha]\times\{i\}}\right) \qquad \text{for } i \in [\ell]. \tag{3}$$

Given access to k nodes, we regard their information as a $k \times \alpha \times \ell$ tensor and apply the classical reconstruction algorithm to the ith slice for $i \in [\ell]$. Since this slice is in fact $\textsc{Encode}\left(x_i^{(1)}, x_i^{(2)}, \ldots, x_i^{(B)}\right)$, we retrieve the data units $\left(x_i^{(1)}, x_i^{(2)}, \ldots, x_i^{(B)}\right)$, which correspond to the ith coordinates of the data blocks $\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \ldots, \mathbf{x}^{(B)}$. Hence, the algorithm given by (2) retrieves the complete collection of data blocks of the B users. A similar setup holds for the repair algorithm. As a result, $\textsc{Encode}(\ell)$, $\textsc{Reconstruct}(T;\ell)$ and $\textsc{Repair}(t, T;\ell)$ describe an $(n, k, d, \alpha\ell, B\ell)$ DSS code.

Suppose after the edits, the data blocks are updated to $\tilde{\mathbf{x}}^{(1)}, \tilde{\mathbf{x}}^{(2)}, \ldots, \tilde{\mathbf{x}}^{(B)}$, each of length ℓ'. One straightforward approach to maintain reconstruction and repair properties after the updates is to ensure that the information stored at the n nodes is given by $\textsc{Encode}(\ell')\left(\tilde{\mathbf{x}}^{(1)}, \tilde{\mathbf{x}}^{(2)}, \ldots, \tilde{\mathbf{x}}^{(B)}\right)$. Unfortunately, as we will see in the next example, such an approach requires a communication cost of at least $(\ell - 1)\log q$ bits in the worst case scenario.

Example 1 (Generalized). Recall the systematic single parity $[3, 2]$ MDS code and let the data blocks of users 1 and 2 be $\mathbf{u}^{(1)}, \mathbf{u}^{(2)} \in \mathbb{F}_q^{\ell}$, respectively. For this code, we may use (1) and rewrite the information at the storage nodes as

$$\textsc{Encode}(\ell)\left(\mathbf{u}^{(1)}, \mathbf{u}^{(2)}\right) = \begin{pmatrix} \mathbf{u}^{(1)} \\ \mathbf{u}^{(2)} \\ \mathbf{u}^{(1)} + \mathbf{u}^{(2)} \end{pmatrix}. \tag{4}$$

Recall that in the case for MDS codes we have $\alpha = 1$. Hence, the image of the encoding map is simply written as a matrix, instead of a three dimensional tensor. This notational convention applies to all subsequent examples in the paper.

[3]For compactness, unless stated otherwise, we drop the transposition symbol T when writing the $B \times \ell$ array representing the B data blocks of length ℓ. In other words, we write $\left(\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \ldots, \mathbf{x}^{(B)}\right)$ instead of $\left(\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \ldots, \mathbf{x}^{(B)}\right)^T$.

Proposition 1. *Let* $\mathbf{u}^{(1)}$ *and* $\mathbf{u}^{(2)}$ *be strings of length* ℓ*. Assume that exactly one deletion has occurred in* $\mathbf{u}^{(1)}$ *and* $\mathbf{u}^{(2)}$*. Let* $\tilde{\mathbf{u}}^{(1)}$ *and* $\tilde{\mathbf{u}}^{(2)}$ *denote the respective edited strings, and suppose that the information to be updated at the storage nodes is given by* (4)*. Then, in the worst case over all possible edit locations, the total communication cost is at least* $\ell - 1$ *symbols, or* $(\ell - 1)\log q$ *bits, independent of the network topology between users and storage nodes.*

Proof. To prove the claimed result, we adapt the *fooling set method* from communication complexity (see [37, §13.2.1]), standardly used to lower bound the deterministic communication complexity of a function.

Assume that $u_\ell^{(1)}$ and $u_1^{(2)}$ are the deleted coordinates. Then the task of node 3 is to update its value to $\left(u_1^{(1)} + u_2^{(2)}, u_2^{(1)} + u_3^{(2)}, \ldots, u_{\ell-1}^{(1)} + u_\ell^{(2)}\right)$.

Let $f : \mathbb{F}_q^\ell \times \mathbb{F}_q^\ell \to \mathbb{F}_q^{\ell-1}$ be a function with $f(\mathbf{u}^{(1)}, \mathbf{u}^{(2)}) = \left(u_1^{(1)} + u_2^{(2)}, u_2^{(1)} + u_3^{(2)}, \ldots, u_{\ell-1}^{(1)} + u_\ell^{(2)}\right)$. A *fooling set* for f of size M is a subset $S \subseteq \mathbb{F}_q^\ell \times \mathbb{F}_q^\ell$ and and a value $\mathbf{c} \in \mathbb{F}_q^{\ell-1}$ such that (a) for every $\left(\mathbf{u}^{(1)}, \mathbf{u}^{(2)}\right) \in S$, $f\left(\mathbf{u}^{(1)}, \mathbf{u}^{(2)}\right) = \mathbf{c}$ and (b) for every distinct $\left(\mathbf{u}^{(1)}, \mathbf{u}^{(2)}\right), \left(\mathbf{v}^{(1)}, \mathbf{v}^{(2)}\right) \in S$, either $f(\mathbf{u}^{(1)}, \mathbf{v}^{(2)}) \neq \mathbf{c}$ or $f(\mathbf{v}^{(1)}, \mathbf{u}^{(2)}) \neq \mathbf{c}$. One can show that if f has a fooling set of size M, then the total deterministic communication cost for any protocol computing f is at least $\log M$.

To prove the claim, we exhibit next a fooling set of size $q^{\ell-1}$ for the function of interest. Consider the subset of $\mathbb{F}_q^\ell \times \mathbb{F}_q^\ell$ of size $q^{\ell-1}$ defined as

$$S = \left\{ \left(\mathbf{u}^{(1)}, \mathbf{u}^{(2)}\right) : \mathbf{u}^{(1)} = \left(0, u_2^{(1)}, u_3^{(1)}, \ldots, u_\ell^{(1)}\right), \mathbf{u}^{(2)} = \left(-u_2^{(1)}, -u_3^{(1)}, \ldots, -u_\ell^{(1)}, 0\right) \right\},$$

and let $\mathbf{c} = \mathbf{0}$. Then $f\left(\mathbf{u}^{(1)}, \mathbf{u}^{(2)}\right) = \mathbf{0}$ for all $(\mathbf{u}^{(1)}, \mathbf{u}^{(2)}) \in S$. Furthermore, if $\mathbf{u}^{(1)} \neq \mathbf{v}^{(1)}$, or equivalently, if $\left(0, u_2^{(1)}, u_3^{(1)}, \ldots, u_\ell^{(1)}\right) \neq \left(0, v_2^{(1)}, v_3^{(1)}, \ldots, v_\ell^{(1)}\right)$, then we can check that $f\left(\mathbf{u}^{(1)}, \mathbf{v}^{(2)}\right) = \left(u_2^{(1)} - v_2^{(1)}, u_3^{(1)} - v_3^{(1)}, \ldots, u_\ell^{(1)} - v_\ell^{(1)}\right) \neq \mathbf{0}$. Therefore, S is a fooling set of size $q^{\ell-1}$. □

For the general case involving more than two users and more than three storage nodes, one can focus on the worst case scenario in which two users each have a single deletion and need to update a parity check value in a common, connected node. The proof of Proposition 1 can be easily modified to show that in this case, the worst case communication cost remains $(\ell - 1)$ symbols, or $(\ell - 1)\log q$ bits.

3.1 Update Protocols: Beyond the Worst Case

In what follows, we describe a straightforward update protocol for the traditional encoding scheme with edits that are not necessarily confined to the worst case configuration.

Consider a $(n, k, d, \alpha\ell, B\ell)$ DSS code with the data blocks of the users equal to $\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \ldots, \mathbf{x}^{(B)}$, and the information stored in the storage nodes equal to $\textsc{Encode}(\ell)\left(\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \ldots, \mathbf{x}^{(B)}\right)$. We consider the uniform edits model and for simplicity, assume that there is a single deletion[4] at coordinate i_s in data block $\mathbf{x}^{(s)}$, for $s \in [B]$. Hence, the updated data block length equals $\ell' = \ell - 1$. Let $\tilde{\mathbf{x}}^{(1)}, \tilde{\mathbf{x}}^{(2)}, \ldots, \tilde{\mathbf{x}}^{(B)}$ be the edited data blocks. In the traditional encoding scheme, to preserve reconstruction and repair properties, we require the information stored by the nodes to be updated to $\textsc{Encode}(\ell - 1)\left(\tilde{\mathbf{x}}^{(1)}, \tilde{\mathbf{x}}^{(2)}, \ldots, \tilde{\mathbf{x}}^{(B)}\right)$.

[4]Insertions can be treated in an almost identical manner as deletions and will not be explicitly discussed.

Example 1 demonstrates that in the worst case, a user needs to transmit $(\ell-1)\log q$ symbols to its connected nodes. Clearly, for most edit scenarios, certain portions of the encoded data do not need to be updated. Continuing Example 1, suppose $\tilde{\mathbf{x}} = (x_1, x_2, \ldots, x_{\ell-2}, x_\ell)$ with deletion at coordinate $\ell-1$ and $\tilde{\mathbf{y}} = (y_1, y_2, \ldots, y_{\ell-2}, y_{\ell-1})$ with deletion at coordinate ℓ. Node 3 needs to update its string to $\tilde{\mathbf{x}} + \tilde{\mathbf{y}} = (x_1+y_1, x_2+y_2, \ldots, x_{\ell-2}+y_{\ell-2}, x_\ell + y_{\ell-1})$, and it suffices for users 1 and 2 to transmit x_ℓ and $y_{\ell-1}$, respectively, to node 3. Indeed, to compute $\tilde{\mathbf{x}} + \tilde{\mathbf{y}}$, node 3 only needs to compute the last coordinate as the other coordinates remain the same.

Suppose that all users send their deleted coordinates to a designated central storage node, and that the designated central node computes $i_{\max} \triangleq \max_{s\in[B]} i_s$ and $i_{\min} \triangleq \min_{s\in[B]} i_s$.

Define $I = \{i \in [\ell] : i_{\min} \le i \le i_{\max} - 1\}$. Since $\left(\tilde{x}_i^{(1)}, \tilde{x}_i^{(2)}, \ldots, \tilde{x}_i^{(B)}\right) = \left(x_i^{(1)}, x_i^{(2)}, \ldots, x_i^{(B)}\right)$ for $i < i_{\min}$, and $\left(\tilde{x}_i^{(1)}, \tilde{x}_i^{(2)}, \ldots, \tilde{x}_i^{(B)}\right) = \left(x_{i+1}^{(1)}, x_{i+1}^{(2)}, \ldots, x_{i+1}^{(B)}\right)$ for $i \ge i_{\max}$, we have

$$\textsc{Encode}(\ell-1)\left(\tilde{\mathbf{x}}^{(1)}, \tilde{\mathbf{x}}^{(2)}, \ldots, \tilde{\mathbf{x}}^{(B)}\right)\Big|_{[n]\times[\alpha]\times\{i\}} = \begin{cases} \textsc{Encode}(\ell)\left(\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \ldots, \mathbf{x}^{(B)}\right)\big|_{[n]\times[\alpha]\times\{i\}}, & \text{if } i < i_{\min}, \\ \textsc{Encode}(\ell)\left(\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \ldots, \mathbf{x}^{(B)}\right)\big|_{[n]\times[\alpha]\times\{i+1\}}, & \text{if } i \ge i_{\max}. \end{cases} \tag{5}$$

In other words, the information stored at coordinates in $[n]\times[\alpha]\times([\ell'] \setminus I)$ need not be updated. Therefore, it suffices to compute $\textsc{Encode}(\ell-1)\left(\tilde{\mathbf{x}}^{(1)}, \tilde{\mathbf{x}}^{(2)}, \ldots, \tilde{\mathbf{x}}^{(B)}\right)\big|_{[n]\times[\alpha]\times I}$ as given by

$$\textsc{Encode}(\ell-1)\left(\tilde{\mathbf{x}}^{(1)}, \tilde{\mathbf{x}}^{(2)}, \ldots, \tilde{\mathbf{x}}^{(B)}\right)\Big|_{[n]\times[\alpha]\times I} = \textsc{Encode}(|I|)\left(\tilde{\mathbf{x}}_I^{(1)}, \tilde{\mathbf{x}}_I^{(2)}, \ldots, \tilde{\mathbf{x}}_I^{(B)}\right), \tag{6}$$

where $\tilde{\mathbf{x}}_I^{(s)}$ is the updated string restricted to the coordinates in I. In other words, $\tilde{\mathbf{x}}_I^{(s)} \triangleq \left(x_{i_{\min}}^{(s)}, x_{i_{\min}+1}^{(s)}, \ldots, x_{i_s-1}^{(s)}, x_{i_s+1}^{(s)}, \ldots, x_{i_{\max}}^{(s)}\right)$.

Consequently, equipped with $\tilde{\mathbf{x}}_I^{(s)}$ for $s \in [B]$ and equations (5) and (6), the nodes are able to compute $\textsc{Encode}(\ell-1)\left(\tilde{\mathbf{x}}^{(1)}, \tilde{\mathbf{x}}^{(2)}, \ldots, \tilde{\mathbf{x}}^{(B)}\right)$. We call this update protocol *Scheme T*.

Proposition 2 (Scheme T)**.** *Consider an $(n, k, d, \alpha\ell, B\ell)$ DSS code and assume single deletions in the user data blocks. The updates in accordance to Scheme T result in an $(n, k, d, \alpha(\ell-1), B(\ell-1))$ DSS code. Each user needs to send out $|I|\log q$ bits to a connected storage node, which in the worst case equals $(\ell-1)\log q$ bits. Here, $I = \{i \in [\ell] : i_{\min} \le i \le i_{\max} - 1\}$, where $i_{\max} = \max_{s\in[B]} i_s$ and $i_{\min} = \min_{s\in[B]} i_s$.*

In Appendix B, we compute the expected communication cost between a user and a connected storage node for a number of probabilistic edit models, and show that for all models considered this cost is of the order of $\ell \log q$ bits, i.e., of the order of the worst case communication scenario rate. The analysis is based on order statistics that characterizes the average span of deletions in the data blocks.

4 Synchronization Schemes with Order-Optimal Communication Cost

To avoid repeated transmissions of all data blocks of users, one needs to develop encoding methods that work around the problems associated with the simple protocol T. We present next our main

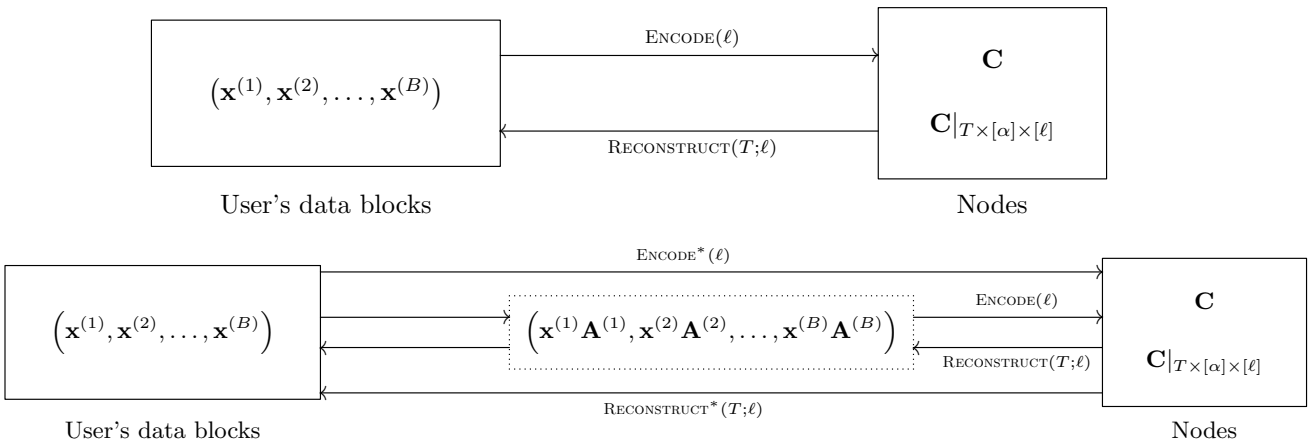

Figure 1: a) A diagram illustrating how traditional encoding extends an existing (n, k, d, α, B) DSS code to an $(n, k, d, \alpha\ell, B\ell)$ DSS code so as to handle data blocks of length ℓ. b) Intermediary encoding that retains both reconstruction and repair functionalities.

results, a collection of synchronization schemes that achieve communication cost of $O(\log \ell + \log q)$ bits as opposed to $O(\ell \log q)$ required in the T protocol setting. Our schemes are based on the following observation, apparent from Example 1: if one insists on using (1) to encode data, then edit updates inevitably incur a communication cost of the order of the data block lengths. Hence, the idea is to introduce an intermediary encoding algorithm that preserves the reconstruction and repair properties of the DSS code, while allowing for significantly lower edit information communication costs. This intermediary encoding algorithm is described in what follows and is illustrated in Figure 1.

4.1 Intermediary Encoding

Let $\left(\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \ldots, \mathbf{x}^{(B)}\right) \in \mathbb{F}_q^{B\times\ell}$ and let $\mathbf{A}^{(1)}, \mathbf{A}^{(2)}, \ldots, \mathbf{A}^{(B)}$ be invertible $\ell \times \ell$ matrices over $\mathbb{F}_q$. We define ENCODE$^* : \mathbb{F}_q^{B\times\ell} \to \mathbb{F}_q^{n\times\alpha\times\ell}$ as

$$\text{ENCODE}^*(\ell)\left(\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \ldots, \mathbf{x}^{(B)}\right) \triangleq \text{ENCODE}(\ell)\left(\mathbf{x}^{(1)}\mathbf{A}^{(1)}, \mathbf{x}^{(2)}\mathbf{A}^{(2)}, \ldots, \mathbf{x}^{(B)}\mathbf{A}^{(B)}\right). \tag{1*}$$

Next, we show that this encoding function preserves reconstruction and repair capabilities. The proof for the repair capability is straightforward. For $t \in [n]$ and $T \in \binom{[n]\setminus\{t\}}{d}$, simply define REPAIR$^*(t, T; \ell) : \mathbb{F}_q^{d\times\alpha\times\ell} \to \mathbb{F}_q^{\alpha\times\ell}$ via

$$\text{REPAIR}^*(t, T; \ell)(\mathbf{C}|_{T\times[\alpha]\times[\ell]}) \triangleq \text{REPAIR}(t, T; \ell)(\mathbf{C}|_{T\times[\alpha]\times[\ell]}) \tag{3*}$$

Since the right hand side equals $\mathbf{C}|_{\{t\}\times[\alpha]\times[\ell]}$, the repair property holds.

For the reconstruction property, consider the following algorithm for any $T \in \binom{[n]}{k}$.

Algorithm 2$*$: $\textsc{Reconstruct}^*(T;\ell)(\mathbf{C}|_{T\times[\alpha]\times[\ell]})$ $(2*)$

$\left(\mathbf{y}^{(1)},\mathbf{y}^{(2)},\ldots,\mathbf{y}^{(B)}\right) \leftarrow \textsc{Reconstruct}(T;\ell)(\mathbf{C}|_{T\times[\alpha]\times[\ell]})$
for $s \in [B]$ **do**
  $\hat{\mathbf{x}}^{(s)} \leftarrow \mathbf{y}^{(s)}\left(\mathbf{A}^{(s)}\right)^{-1}$
end
return $\left(\hat{\mathbf{x}}^{(1)},\hat{\mathbf{x}}^{(2)},\ldots,\hat{\mathbf{x}}^{(B)}\right)$

The reconstruction algorithm produces the correct output since from $\mathbf{y}^{(s)} = \mathbf{x}^{(s)}\mathbf{A}^{(s)}$ for $s \in [B]$, we have $\hat{\mathbf{x}}^{(s)} = \mathbf{x}^{(s)}$.

We summarize the observations in the following proposition.

Proposition 3. *Consider an (n,k,d,α,B) DSS code with functions* Encode, Reconstruct *and* Repair. *For any positive integer ℓ and B invertible $\ell\times\ell$ matrices $\mathbf{A}^{(1)},\mathbf{A}^{(2)},\ldots,\mathbf{A}^{(B)}$, the functions* $\textsc{Encode}^*$, $\textsc{Reconstruct}^*$ *and* $\textsc{Repair}^*$ *given by* $(1*)$, $(2*)$ *and* $(3*)$ *describe an $(n,k,d,\alpha\ell,B\ell)$ DSS code.*

Simple modifications of $(1*)$ that allow for (a) systematic DSS coding and for (b) variable data block lengths are described in what follows.

(a) **Systematic DSS**. A *systematic* DSS code is a DSS code with the property that the B data blocks are explicitly stored amongst a set of k nodes, termed the systematic nodes. In other words, the content of the systematic nodes are the blocks $\mathbf{x}^{(s)}$.

Assume a systematic DSS code with functionalities Encode, Reconstruct, Repair. As before, let $\mathbf{A}^{(1)},\mathbf{A}^{(2)},\ldots,\mathbf{A}^{(B)}$ be invertible $\ell\times\ell$ matrices and define a systematic encoding function $\textsc{Encode}^*(\ell)\left(\mathbf{x}^{(1)},\mathbf{x}^{(2)},\ldots,\mathbf{x}^{(B)}\right)$ according to

$$\textsc{Encode}^*(\ell)\left(\mathbf{x}^{(1)},\mathbf{x}^{(2)},\ldots,\mathbf{x}^{(B)}\right)\Big|_{\{t\}\times[\alpha]\times[\ell]}$$
$$\triangleq \begin{cases} \textsc{Encode}(\ell)\left(\mathbf{x}^{(1)},\mathbf{x}^{(2)},\ldots,\mathbf{x}^{(B)}\right)\big|_{\{t\}\times[\alpha]\times[\ell]}, & \text{if } t \text{ is a systematic node,} \\ \textsc{Encode}(\ell)\left(\mathbf{x}^{(1)}\mathbf{A}^{(1)},\mathbf{x}^{(2)}\mathbf{A}^{(2)},\ldots,\mathbf{x}^{(B)}\mathbf{A}^{(B)}\right)\big|_{\{t\}\times[\alpha]\times[\ell]}, & \text{otherwise.} \end{cases}$$

Let $\mathbf{C}$ be the resulting information stored over the n nodes. The reconstruction and repair algorithms may be used unaltered provided that simple pre-processing of $\mathbf{C}|_{T\times[\alpha]\times[\ell]}$ is performed first. Specifically, if $t \in T$ and t is a systematic node, we first let $\mathbf{C}|_{\{t\}\times[\alpha]\times[\ell]} = \left(\mathbf{x}^{(s_1)},\mathbf{x}^{(s_2)},\ldots,\mathbf{x}^{(s_\alpha)}\right)$. Then, we modify this information to

$$\left(\mathbf{x}^{(s_1)}\mathbf{A}^{(s_1)},\mathbf{x}^{(s_2)}\mathbf{A}^{(s_2)},\ldots,\mathbf{x}^{(s_\alpha)}\mathbf{A}^{(s_\alpha)}\right).$$

If $\mathbf{C}'$ denotes the tensor resulting from the previous computation, reconstruction and repair are performed according to $\textsc{Reconstruct}^*(T;\ell)(\mathbf{C}')$ and $\textsc{Repair}^*(T;\ell)(\mathbf{C}')$, respectively.

(b) **Data blocks of variable lengths**. For $s \in [B]$, let $\mathbf{x}^{(s)}$ be of length ℓ_s and let $\mathbf{A}^{(s)}$ be a *right* invertible $\ell_s\times\ell$ matrix, where $\ell = \max_{s\in[B]}\ell_s$. Then we define $\textsc{Encode}^*(\ell) : \mathbb{F}_q^{\ell_1}\times\mathbb{F}_q^{\ell_2}\times\cdots\times\mathbb{F}_q^{\ell_B} \to \mathbb{F}_q^{n\times\alpha\times\ell}$ via $(1*)$. If we further define the reconstruction and repair functions via

(2∗) and (3∗), we arrive at an $(n, k, d, \alpha\ell, B\ell)$ DSS code as before. The right inverse of $\mathbf{A}^{(s)}$ is required for reconstruction and for simplicity, we use the notation $\left(\mathbf{A}^{(s)}\right)^{-1}$ for both the inverse and right inverse of a square and rectangular matrix, respectively.

Note that when $\mathbf{A}^{(s)} = \mathbf{I}$ for all $s \in [B]$, we recover the traditional encoding function. In the next subsection, we describe how to choose the matrices $\mathbf{A}^{(s)}$ and accompanying update protocols so as to ensure significantly lower communication cost between a user and connected storage node as compared to the traditional scheme.

4.2 Synchronization Schemes

Suppose the data blocks are edited and the resulting data blocks are of length ℓ'. The idea behind our approach is to request the users to modify their respective matrices $\mathbf{A}^{(s)}$ to be invertible $\ell' \times \ell'$ matrices $\tilde{\mathbf{A}}^{(s)}$ according to the edits made. Then, users can only transmit the *locations* and *values* of their edits rather than a whole span of values, with the storage nodes still being able to update their respective information so that (1∗) holds. Since the matrices $\tilde{\mathbf{A}}^{(s)}$ are designed to be invertible, the resulting system remains an $(n, k, d, \alpha\ell', B\ell')$ DSS code.

We propose three different update schemes for the matrices $\tilde{\mathbf{A}}^{(s)}$ based on the frequency and extent to which edits are made, catering to the need of [14]:

(4.1) **Semi-static Data**. Here, we assume that only a *constant fraction* of the data blocks is edited by users so that most data blocks retained their original length ℓ. In this case, the matrices $\tilde{\mathbf{A}}^{(s)}$ – albeit modified – remain of dimension $\ell \times \ell$. The most appropriate choice for the matrices are permutation matrices, i.e., 0-1 matrices with exactly one 1 per row and per column. We observe that each node stores $\alpha\ell$ symbols.

(4.2) **Hot Data**. In contrast to the semi-static case, one may also assume that a significant proportion of the data blocks are edited by users, as is the case for hot data[5]. In this case, suppose that the resulting data blocks are of length $\ell' < \ell$; then the storage requirements reduce to each node storing $\alpha\ell' < \alpha\ell$ symbols. The matrices $\tilde{\mathbf{A}}^{(s)}$ have dimension $\ell' \times \ell'$ and an appropriate choice for them are Vandermonde and Cauchy matrices (see [33]), to be discussed in more detail throughout the next sections.

Scheme Based on Permutation Matrices. Although this scheme applies for general nonuniform edit models, and consequently to data blocks of variables lengths, for simplicity we assume edits of the form of a single deletion or insertion. As before, in the nonuniform setting we pad the shorter data blocks by an appropriate number of zeros. To do so, we represent the file by padding the data string with zeroes. So, a deletion in $\mathbf{x}$ at position 2 results in $\tilde{\mathbf{x}} = (x_1, x_3, \ldots, x_{\ell-1}, 0)$.

Let data block $\mathbf{x}^{(s)}$ be edited at coordinate i_s. Recall that we associate with $\mathbf{x}^{(s)}$ an $\ell \times \ell$ matrix $\mathbf{A}^{(s)}$. The matrix $\mathbf{A}^{(s)}$ is initialized to the identity matrix $\mathbf{I}$ and it remains a permutation matrix after each update. Roughly speaking, the storage nodes maintain the coded information in the original order. Since with each edit this order changes, the permutation matrix $\mathbf{A}^{(s)}$ is used to keep track of the order in the data blocks relative to that in the storage nodes. Hence, $\mathbf{A}^{(s)}$ indicates

[5]In many application, hot data is left uncoded in order to facilitate quick access to information and eliminate the need for re-encoding. Our scheme mitigates both the issues of access and re-encoding, while allowing higher level of data integrity through distributed coding.

that instead of editing "position i_s" of the check nodes, one has to edit a position "position j_s" in the original order. These assertions are stated and proved rigorously in the proof of Proposition 4.

Since permutation matrices are clearly invertible, the approach results in an $(n, k, d, \alpha\ell, B\ell)$ DSS code by Proposition 3. The permutation matrix intermediary encoding scheme is detailed in Scheme P.

Scheme P: Suppose $\mathbf{x}^{(s)}$ is edited at coordinate i_s.

1 **if** edit is a deletion **then**
2 | $j_s \leftarrow$ coordinate where the i_sth row $\mathbf{A}^{(s)}$ is one (note: $\mathbf{A}^{(s)}$ is a permutation matrix)
3 | shift i_sth row of $\mathbf{A}^{(s)}$ to the last row
4 **else**
5 | $j_s \leftarrow$ coordinate where the last row $\mathbf{A}^{(s)}$ is one
6 | shift last row of $\mathbf{A}^{(s)}$ to the i_sth row
7 **end**
8 User s sends to the connected storage nodes $N(s)$: the value affected, using x ($\log q$ bits), the type of edit – insertion or deletion (one bit), and the coordinate j_s ($\log \ell$ bits)
9 **for** $t \in N(s)$ **do**
10 | Compute $\mathbf{d} = \textsc{Encode}(\ell)\,(\mathbf{0}, \ldots, \mathbf{0}, x\mathbf{e}_{j_s}, \mathbf{0}, \ldots, \mathbf{0})\,|_{\{t\}\times[\alpha]\times[\ell]}$
11 | **if** edit is a deletion **then**
12 | | subtract $\mathbf{d}$ from coordinate j_s at each storage node
13 | **else**
14 | | add $\mathbf{d}$ to coordinate j_s at each storage node
15 | **end**
16 **end**

Proposition 4 (Scheme P)**.** *Consider an $(n, k, d, \alpha\ell, B\ell)$ DSS code and assume a single edit for a single user. The updates in accordance to Scheme P result in an $(n, k, d, \alpha\ell, B\ell)$ DSS code and the user needs to communicate $\log \ell + \log q$ bits to a connected storage node to update her/his information.*

Before we prove this proposition, we illustrate Scheme P via an example.

Example 2. Consider a simple example of two data blocks $\mathbf{u}^{(1)}$ and $\mathbf{u}^{(2)}$, shown below, for which $\ell = 5$, as part of a $[3, 2]$ MDS code over $\mathbb{F}_5$. Let $\mathbf{D}$ denote the tensor given by $\textsc{Encode}(\ell)\,(\mathbf{0}, \ldots, \mathbf{0}, x\mathbf{e}_{j_s}, \mathbf{0}, \ldots, \mathbf{0})$ in line 10.

$\mathbf{u}^{(1)}$	$\mathbf{u}^{(2)}$	edit	$\mathbf{D}$	$\textsc{Encode}(5)\left(\mathbf{u}^{(1)}\mathbf{A}^{(1)}, \mathbf{u}^{(2)}\mathbf{A}^{(2)}\right)$	$\mathbf{A}^{(1)}$	$\mathbf{A}^{(2)}$
$(1,2,3,4,4)$	$(1,1,1,1,1)$	—	—	$\begin{pmatrix} 1 & 2 & 3 & 4 & 4 \\ 1 & 1 & 1 & 1 & 1 \\ 2 & 3 & 4 & 0 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}$
$(1,3,4,4,0)$	$(1,1,1,1,1)$	deletion at position 2 of $\mathbf{x}$	$\begin{pmatrix} 0 & 2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 3 & 4 & 4 \\ 1 & 1 & 1 & 1 & 1 \\ 2 & 1 & 4 & 0 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}$
$(1,3,4,0,0)$	$(1,1,1,1,1)$	deletion at position 3 of $\mathbf{x}$	$\begin{pmatrix} 0 & 0 & 0 & 4 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 4 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 3 & 0 & 4 \\ 1 & 1 & 1 & 1 & 1 \\ 2 & 1 & 4 & 1 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}$
$(1,4,3,4,0)$	$(1,1,1,1,1)$	insertion of 4 at position 2 of $\mathbf{x}$	$\begin{pmatrix} 0 & 0 & 0 & 4 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 4 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 3 & 4 & 4 \\ 1 & 1 & 1 & 1 & 1 \\ 2 & 1 & 4 & 0 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}$

We reconstruct the data blocks using nodes or rows 1 and 3.

1. From row 1 and 3, we infer that row 2 is $(1,1,1,1,1)$.
2. Hence, we obtain

$$\hat{\mathbf{u}}^{(1)} = (1,0,3,4,4)\left(\mathbf{A}^{(1)}\right)^{-1} = (1,4,3,4,0),$$
$$\hat{\mathbf{u}}^{(2)} = (1,1,1,1,1)\left(\mathbf{A}^{(2)}\right)^{-1} = (1,1,1,1,1),$$

as desired.

Proof of Proposition 4. We demonstrate the correctness of Scheme P for the case of a deletion. Insertions are handled similarly.

We first show that the updates performed as part of the scheme result in an $(n,k,d,\alpha\ell,B\ell)$ DSS code. It suffices to show that the nodes store $\textsc{Encode}(\ell)\left(\mathbf{x}^{(1)}\mathbf{A}^{(1)}, \mathbf{x}^{(2)}\mathbf{A}^{(2)}, \ldots, \tilde{\mathbf{x}}^{(s)}\tilde{\mathbf{A}}^{(s)}, \ldots, \mathbf{x}^{(B)}\mathbf{A}^{(B)}\right)$, where $\tilde{\mathbf{A}}^{(s)}$ is the matrix resulting from instruction performed from line 1 to line 7. Observe that prior to the edit, the nodes stored $\textsc{Encode}(\ell)\left(\mathbf{x}^{(1)}\mathbf{A}^{(1)}, \mathbf{x}^{(2)}\mathbf{A}^{(2)}, \ldots, \mathbf{x}^{(s)}\mathbf{A}^{(s)}, \ldots, \mathbf{x}^{(B)}\mathbf{A}^{(B)}\right)$. Define

$$\begin{aligned}\mathbf{D} &\triangleq \textsc{Encode}(\ell)\left(\mathbf{x}^{(1)}\mathbf{A}^{(1)}, \mathbf{x}^{(2)}\mathbf{A}^{(2)}, \ldots, \mathbf{x}^{(s)}\mathbf{A}^{(s)}, \ldots, \mathbf{x}^{(B)}\mathbf{A}^{(B)}\right) \\ &\quad - \textsc{Encode}(\ell)\left(\mathbf{x}^{(1)}\mathbf{A}^{(1)}, \mathbf{x}^{(2)}\mathbf{A}^{(2)}, \ldots, \tilde{\mathbf{x}}^{(s)}\tilde{\mathbf{A}}^{(s)}, \ldots, \mathbf{x}^{(B)}\mathbf{A}^{(B)}\right) \\ &= \textsc{Encode}(\ell)\left(\mathbf{0}, \ldots, \mathbf{0}, \mathbf{x}^{(s)}\mathbf{A}^{(s)} - \tilde{\mathbf{x}}^{(s)}\tilde{\mathbf{A}}^{(s)}, \mathbf{0}, \ldots, \mathbf{0}\right),\end{aligned}$$

and suppose that the following claim were true

$$\mathbf{x}^{(s)}\mathbf{A}^{(s)} - \tilde{\mathbf{x}}^{(s)}\tilde{\mathbf{A}}^{(s)} = x\mathbf{e}_{j_s}. \tag{7}$$

Then, the updates in lines 9 to 16 have a net effect of subtracting $\mathbf{D}$ from the information stored at the nodes (recall that $\textsc{Encode}(\mathbf{e}_s)|_{\{t\}\times[\alpha]} = \mathbf{0}$ for $t \notin N(s)$). This completes the proof.

Hence, it remains to show that (7) is correct. Without loss of generality, let $\mathbf{x}^{(s)} = (x_1, x_2, \ldots, x_\ell)$. Then $\tilde{\mathbf{x}}^{(s)} = (x_1, x_2, \ldots, x_{i_s-1}, x_{i_s+1}, \ldots, x_\ell, 0)$ and $x = x_{i_s}$.

Since $\mathbf{A}^{(s)}|_{\{i\}\times[\ell]} = \tilde{\mathbf{A}}^{(s)}|_{\{i\}\times[\ell]}$ for $i \leq i_s - 1$, we have $x_i\mathbf{A}^{(s)}|_{\{i\}\times[\ell]} - x_i\tilde{\mathbf{A}}^{(s)}|_{\{i\}\times[\ell]} = \mathbf{0}$. Similarly, $x_i\mathbf{A}^{(s)}|_{\{i\}\times[\ell]} - x_i\tilde{\mathbf{A}}^{(s)}|_{\{i-1\}\times[\ell]} = \mathbf{0}$ for $i_s + 1 \leq i \leq \ell$. Hence, the left hand side yields $x_{i_s}\tilde{\mathbf{A}}^{(s)}|_{\{i_s\}\times[\ell]} - 0 \cdot \tilde{\mathbf{A}}^{(s)}|_{\{\ell\}\times[\ell]} = x\mathbf{e}_{j_s}$, as needed. □

Scheme Based on Vandermonde Matrices. As for the case of coding with permutation matrices, the scheme using Vandermonde matrices applies to nonuniform edits, but for simplicity we assume that there is a single deletion at coordinate i_s in data block $\mathbf{x}^{(s)}$, for $s \in [B]$. Nevertheless, at the end of the section, we briefly outline minor algorithmic changes that need to be performed in order to accommodate the non-uniform deletion model.

When one deletion is present, the updated data block length is $\ell' = \ell - 1$. Recall that an $\ell \times \ell$ Vandermonde matrix is an invertible matrix of the form

$$\begin{bmatrix} a_1 & a_1^2 & \ldots & a_1^{\ell-1} & a_1^\ell \\ a_2 & a_2^2 & \ldots & a_2^{\ell-1} & a_2^\ell \\ \ldots & \ldots & \ldots & \ldots & \ldots \\ a_{\ell-1} & a_{\ell-1}^2 & \ldots & a_{\ell-1}^{\ell-1} & a_{\ell-1}^\ell \\ a_\ell & a_\ell^2 & \ldots & a_\ell^{\ell-1} & a_\ell^\ell \end{bmatrix},$$

where $a_1, \ldots, a_\ell$ are distinct values over an appropriate field.

Recall that we associate with $\mathbf{x}^{(s)}$ an $\ell \times \ell$ matrix $\mathbf{A}^{(s)}$. After synchronization, we want the updated matrix $\tilde{\mathbf{A}}^{(s)}$ to be of dimension $(\ell-1) \times (\ell-1)$ and invertible, and the information in the n storage nodes reduced to $\alpha \times (\ell - 1)$ arrays.

The deleted values $x_{i_1}^{(1)}, x_{i_2}^{(2)}, \ldots, x_{i_B}^{(B)}$ are stored in the storage nodes as

$$\mathbf{D} \triangleq \textsc{Encode}(\ell)\left(x_{i_1}^{(1)}\mathbf{A}|_{\{i_1\}\times[\ell]}, x_{i_2}^{(2)}\mathbf{A}|_{\{i_2\}\times[\ell]}, \ldots, x_{i_B}^{(B)}\mathbf{A}|_{\{i_B\}\times[\ell]}\right). \tag{8}$$

Hence, when given the values $x_{i_s}^{(s)}$ and positions i_s, each node t may subtract the vector $\mathbf{D}|_{\{t\}\times[\alpha]\times[\ell]}$ from its content. To reduce the size of the storage node arrays, we simply remove the coordinates in the set $[n] \times [\alpha] \times \{\ell\}$. Suppose $\tilde{\mathbf{A}}^{(s)}$ is the $(\ell-1) \times (\ell-1)$ matrix obtained from $\mathbf{A}^{(s)}$ by removing the i_sth row and last column. It is easy to check that *a posteriori* the edit, the storage nodes contain the tensor

$$\textsc{Encode}(\ell - 1)\left(\tilde{\mathbf{x}}^{(1)}\tilde{\mathbf{A}}^{(1)}, \tilde{\mathbf{x}}^{(2)}\tilde{\mathbf{A}}^{(2)}, \ldots, \tilde{\mathbf{x}}^{(B)}\tilde{\mathbf{A}}^{(B)}\right). \tag{9}$$

For the system to be an $(n, k, d, \alpha(\ell-1), B(\ell-1))$ DSS code, we require $\tilde{\mathbf{A}}^{(s)}$ to remain invertible. This is clearly true if the matrix $\mathbf{A}^{(s)}$ is Vandermonde. We refer to the method as Scheme V, the

details of which are given in what follows.

Scheme V: Symbol $\mathbf{x}^{(s)}$ is deleted at coordinate i_s, $s \in [B]$.

for $s \in [B]$ **do**
  User s sends to all connected storage nodes its deleted value deleted – $x_{i_s}^{(s)}$ – ($\log q$ bits) as well as the coordinate i_s of the deletion ($\log \ell$ bits)
  $\mathbf{A}^{(s)} \leftarrow \mathbf{A}^{(s)}$ via removal of the i_sth row and last column
end
for $t \in [n]$ **do**
  Using (8), compute and subtract $\mathbf{D}|_{\{t\}\times[\alpha]\times[\ell]}$
  Remove the (t, j, ℓ)th coordinate, for all $j \in [\alpha]$
end

Proposition 5 (Scheme V)**.** *Consider an $(n, k, d, \alpha\ell, B\ell)$ DSS code and assume single deletions in each user data block. The updates in accordance to Scheme V result in an $(n, k, d, \alpha(\ell-1), B(\ell-1))$ DSS code and each user needs to communicate $\log \ell + \log q$ bits to connected storage nodes to update their information.*

We again illustrate the communication scheme via an example.

Example 3. Assume that $\ell = 4$ and that the DSS code is a $[3, 2]$ MDS code over $\mathbb{F}_5$. As before, we choose two data blocks $\mathbf{u}^{(1)}$ and $\mathbf{u}^{(2)}$ as shown below to illustrate the scheme.

$\mathbf{u}^{(1)}$	i_1	$\mathbf{u}^{(2)}$	i_2	$\mathbf{D}$	$\textsc{Encode}(\ell)\left(\mathbf{u}^{(1)}\mathbf{A}^{(1)}, \mathbf{u}^{(2)}\mathbf{A}^{(2)}\right)$	$\mathbf{A}^{(1)}$	$\mathbf{A}^{(2)}$
$(0,1,0,1)$	—	$(1,0,1,0)$	—	—	$\begin{pmatrix} 2 & 1 & 0 & 2 \\ 2 & 4 & 0 & 3 \\ 4 & 0 & 0 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 2 & 4 & 3 \\ 1 & 3 & 4 & 2 \\ 1 & 4 & 1 & 4 \end{pmatrix}$	$\begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 2 & 4 & 3 \\ 1 & 3 & 4 & 2 \\ 1 & 4 & 1 & 4 \end{pmatrix}$
$(0,1,0)$	4	$(0,1,0)$	1	$\begin{pmatrix} 1 & 4 & 1 & 4 \\ 1 & 1 & 1 & 1 \\ 2 & 0 & 2 & 0 \end{pmatrix}$	$\begin{pmatrix} 1 & 2 & 4 & \not{3} \\ 1 & 3 & 4 & \not{2} \\ 2 & 0 & 3 & \not{0} \end{pmatrix}$	$\begin{pmatrix} 1 & 1 & 1 \\ 1 & 2 & 4 \\ 1 & 3 & 4 \end{pmatrix}$	$\begin{pmatrix} 1 & 2 & 4 \\ 1 & 3 & 4 \\ 1 & 4 & 1 \end{pmatrix}$

We reconstruct the data blocks using nodes or rows 1 and 3.

1. From row 1 and 3, we infer that row 2 is $(1, 3, 4)$.
2. Hence, we obtain

$$\hat{\mathbf{u}}^{(1)} = (1,2,4)\left(\mathbf{A}^{(1)}\right)^{-1} = (0,1,0),$$
$$\hat{\mathbf{u}}^{(2)} = (1,3,4)\left(\mathbf{A}^{(2)}\right)^{-1} = (0,1,0),$$

as desired.

Proof of Proposition 5. As before, it suffices to show that the information stored at the nodes is given by (9). Proceeding in a similar fashion as was done for the proof of Proposition 4, and by referring to the linearity of the encoding maps, the proof reduces to showing that for $s \in [B]$,

$$\left(\mathbf{x}^{(s)}\mathbf{A}^{(s)}\right)\Big|_{[\ell-1]} - \tilde{\mathbf{x}}^{(s)}\tilde{\mathbf{A}}^{(s)} = x_{i_s}^{(s)}\mathbf{A}\Big|_{\{i_s\}\times[\ell-1]}. \tag{10}$$

As before, we check that $x_i^{(s)}\mathbf{A}^{(s)}\Big|_{\{i\}\times[\ell-1]} - x_i^{(s)}\tilde{\mathbf{A}}^{(s)}\Big|_{\{i\}\times[\ell-1]} = \mathbf{0}$ for $i \leq i_s - 1$ and $x_i^{(s)}\mathbf{A}^{(s)}\Big|_{\{i\}\times[\ell-1]} - x_i^{(s)}\tilde{\mathbf{A}}^{(s)}\Big|_{\{i-1\}\times[\ell-1]} = \mathbf{0}$ for $i_s + 1 \leq i \leq \ell$. Hence, the remaining term on the left hand side is $x_{i_s}^{(s)}\mathbf{A}\Big|_{\{i_s\}\times[\ell-1]}$, which establishes (10). □

Remark 4.

(i) **Size of ℓ.** Recall that Vandermonde matrices exist whenever $\ell \leq q$, hence this inequality has to hold true for the scheme to be implementable.

(ii) **Cauchy matrices based encoding as a means of reducing communication complexity**. In Scheme V, all users were required to transmit $(\log \ell + \log q)$ bits to describe their edits. We may save $\log q$ bits in communication complexity by having one of the users, say user 1, transmit only its location.

We achieve this by fixing $\mathbf{A}^{(1)} = \mathbf{I}$. Suppose that $\mathbf{x}^{(1)}$ had a deletion at position i_1. To ensure that $\tilde{\mathbf{A}}^{(1)}$ remains invertible, we remove the i_1th row and i_1th column (in line 3). This in turn forces us to delete the i_1th column in all matrices $\mathbf{A}^{(s)}$. Hence, one needs to ensure that all square submatrices of $\mathbf{A}^{(s)}$ are invertible for $s \geq 2$. It is known that $\ell \times \ell$ Cauchy matrices, taking the general form

$$\begin{bmatrix} 1/(a_1-b_1) & 1/(a_1-b_2) & \dots & 1/(a_1-b_{\ell-1}) & 1/(a_1-b_\ell) \\ 1/(a_2-b_1) & 1/(a_2-b_2) & \dots & 1/(a_2-b_{\ell-1}) & 1/(a_2-b_\ell) \\ \dots & \dots & \dots & \dots & \dots \\ 1/(a_{\ell-1}-b_1) & 1/(a_{\ell-1}-b_2) & \dots & 1/(a_{\ell-1}-b_{\ell-1}) & 1/(a_{\ell-1}-b_\ell) \\ 1/(a_\ell-b_1) & 1/(a_\ell-b_2) & \dots & 1/(a_\ell-b_{\ell-1}) & 1/(a_\ell-b_\ell) \end{bmatrix},$$

for distinct values $a_1, \dots, a_\ell, b_1, \dots, b_\ell$, satisfy this requirement and clearly exist whenever $2\ell \leq q$. Hence, for the protocol to be implementable, the alphabet size and length of data blocks need to satisfy the constraint $2\ell \leq q$.

(iii) **Application to data deduplication.** Deduplication broadly refers to the process of removing duplicate copies of data with the objective of saving storage [15]. Scheme V may easily be integrated into a data deduplication process for a DSS so as to remove duplicates not only amongst the users, but also their redundantly encoded information at the storage nodes.

We describe how to accomplish this task for *post-process* deduplication, i.e., deduplication after the users have already written on their disks certain data blocks, say $(f_1, f_2, \dots, f_e) \in \mathbb{F}_q^e$. Deduplication proceeds as follows:

(I) A central node broadcasts to all users and nodes the data $(f_1, f_2, \dots, f_e)$ to be removed.

(II) For $s \in [B]$, user s scans the string $\mathbf{x}^{(s)}$ for the data string $(f_1, f_2, \dots, f_e)$ and identifies positions $i_{s,1}, i_{s,2}, \dots, i_{s,e}$ where the blocks are stored.

(III) User s transmits positions $i_{s,1}, i_{s,2}, \dots, i_{s,e}$ to all connected storage nodes.

(IV) Each storage node and user updates information as requested by Scheme V in e iterations.

Nonuniform Deletion Model. To conclude this subsection, we describe how to make simple modifications to Scheme V so as to adapt it to the general scenario where each user $\mathbf{x}^{(s)}$ has d_s deletions. Define ℓ' as $\max_{s\in[B]} \ell - d_s$, or equivalently, $\ell - \min_{s\in[B]} d_s$. Our goal is to updated the matrix $\tilde{\mathbf{A}}^{(s)}$ so that it has dimension $(\ell - d_s) \times \ell'$, and reduce the content of the n storage nodes to $\alpha \times \ell'$ arrays.

Suppose $\mathbf{x}^{(s)}$ has deleted values $x^{(s)}_{i_{s1}}, x^{(s)}_{i_{s2}}, \ldots, x^{(s)}_{i_{sd_s}}$. Define $\mathbf{d}^{(s)} \triangleq \sum_{t=1}^{d_s} x^{(s)}_{i_{st}} \mathbf{A}|_{\{i_{st}\}\times[\ell]}$ and modify (8) as

$$\mathbf{D} \triangleq \textsc{Encode}(\ell) \left(\mathbf{d}^{(1)}, \mathbf{d}^{(2)}, \ldots, \mathbf{d}^{(B)}\right). \tag{8'}$$

As before, when given the deleted values and their coordinates, each node t subtracts the vector $\mathbf{D}|_{\{t\}\times[\alpha]\times[\ell]}$. To reduce the size of the storage nodes, we simply remove the coordinates belonging to $[n] \times [\alpha] \times \{i\}$ for $\ell' + 1 \leq i \leq \ell$. Suppose $\tilde{\mathbf{A}}^{(s)}$ is the $(\ell - d_s) \times \ell'$ matrix with the rows corresponding to the d_s deletions and last $\ell - \ell'$ columns removed from $\mathbf{A}^{(s)}$. Then it is not difficult to check that the matrices $\tilde{\mathbf{A}}^{(s)}$ remain Vandermonde if the $\mathbf{A}^{(s)}$ matrices were initially chose to be Vandermonde.

Next, we establish some fundamental limits of communication complexity for the given synchronization of DSS codes under edit protocols, and compare the performance of schemes P, V, and T with these limits.

5 Fundamental Limits and a Tradeoff between Communication Cost and Storage Overhead

Suppose a user's data block is subjected to a *single* edit. Then the communication cost of both Scheme P and Scheme V is approximately $\log \ell + \log q$, for each pair of user and his/her connected storage node. Hence, for a single edit, the communication cost of both schemes is near-optimal.

However, when a data block has an arbitrary number of edits, say d, then the schemes require $d(\log \ell + \log q)$ bits to be communicated and it is unclear if this quantity is optimal, order optimal (i.e., of the same order as the optimal solution) or suboptimal. In the next Section 5.1, we establish a simple lower bound on the communication cost using results of Levenshtein [38] and show that Scheme P and Scheme V are within a constant factor away from the lower bound when $d = o(\ell^{1-\epsilon})$ for a constant $0 < \epsilon < 1$. Hence, under these conditions for d, the schemes are order optimal.

Nevertheless, it is worth pointing out that Scheme P and Scheme V outperform Scheme T for any number of edits. To achieve the communication cost of $O(\log \ell + \log q)$ bits, Scheme P and Scheme V have to store certain structural information regarding the matrices $\mathbf{A}^{(s)}$. As Scheme T does not require this storage overhead, we also analyze the tradeoff between communication cost and storage overhead in Section 5.2. Our findings suggest that the use of Scheme P and Scheme V is preferred to the use of scheme T when the number of edits is $d = o(\ell/\log \ell)$.

5.1 Fundamental Limits

We provide next a lower bound on the communication cost between a user and a connected storage node and show that Scheme P and Scheme V are within a constant factor away from this lower bound provided the number of edits is constant.

Recall that a subsequence of a sequence is itself a sequence that can be obtained from the original sequence by deleting some elements without changing the order of the nonedited elements.

Similarly, a supersequence of a sequence is itself a sequence that can be obtained from the original sequence by inserting some elements without changing the order of the nonedited elements.

Consider the following quantity that counts the number of possible subsequences or supersequences resulting from d edits on a data block of length ℓ, respectively, defined via

$$V(\ell, d) \triangleq \begin{cases} \max_{\mathbf{x}\in\mathbb{F}_q^\ell} |\{\tilde{\mathbf{x}} : \tilde{\mathbf{x}} \text{ result from at most } d \text{ deletions from } \mathbf{x}\}|, \\ \max_{\mathbf{x}\in\mathbb{F}_q^\ell} |\{\tilde{\mathbf{x}} : \tilde{\mathbf{x}} \text{ result from at most } d \text{ insertions and deletions from } \mathbf{x}\}|. \end{cases} \tag{11}$$

We require $\log V(\ell, d)$ bits to describe d deletions (or d insertions/deletions) to a node that contains the original sequence. This shows that each user needs to communicate to a connected storage node at least $\log V(\ell, d)$ bits, given that the storage nodes contain only coded information about the original sequence and that the user knows the positions of the edits. A similar argument for establishing a lower bound for the classical two-way communication protocol was used in [39].

The fundamental combinatorial entity $V(\ell, d)$ was introduced by Levenshtein [40] and has since been studied by a number of authors [38, 40–44]. In particular, it is known (see [38, Eq. (11) and (24)]) that

$$\log V(\ell, d) \geq \begin{cases} \log \binom{\ell-d}{d}, & \text{for deletions only,} \\ \log(q-1)^d \binom{\ell+d}{d}, & \text{for deletions and insertions.} \end{cases}$$

When $d = o(\ell^{1-\epsilon})$ for a constant $0 < \epsilon < 1$, in the asymptotic parameter regime we have $\log \binom{\ell-d}{d} = \Omega(\log \ell)$ and $\log(q-1)^d \binom{\ell+d}{d} = \Omega(\log \ell + \log q)$. Therefore, Scheme P and Scheme V are within a constant factor away from this lower bound.

5.2 Tradeoff between Communication Cost and Storage Overhead

Suppose that during each round of update, a data block has a single edit. In this case, the communication cost between each user and each connected storage node for both Scheme P and Scheme V is approximately $\log \ell + \log q$. However, to achieve this cost, each user needs to store the associated matrix $\mathbf{A}^{(\cdot)}$. Hence, each user a priori requires $\ell^2 \log q$ bits to store this matrix.

However, this stringent storage requirement may be easily relaxed. For any synchronization scheme, we can first store the description of the initial matrices[6]. Subsequently, to generate the matrices $\mathbf{A}^{(s)}$, it suffices for the users to store the modifications to the initial matrices after the edits and we term this information the *storage overhead per edit*. For Scheme P and Scheme V, this overhead amounts to the space allocated for storing the locations of edits; hence, the storage overhead per edit is $\log \ell$ bits. In contrast, for Scheme T, the storage overhead per edit is zero as the matrices $\mathbf{A}^{(\cdot)}$ are always the identity.

We summarize the communication cost and storage overhead features of the various schemes in Table 2.

Storage overhead versus information storage. Consider a single data block that has d edits. Then both Scheme P and Scheme V incur a total storage overhead of $d \log \ell$ bits for the user. But the data block itself is of size $\ell \log q$ bits. Therefore, for desirable storage allocation properties, one would want $d \log \ell = o(\ell)$ or the number of edits to be $o(\ell / \log \ell)$. This implies that Schemes P and V should be used only in the small/moderate edit regime. Nevertheless, it is preferable to use Schemes P and V for semi-static and hot data, respectively, since the former effectively maintains

[6]The initial matrix is the identity matrix for Scheme P and Scheme T, and a Vandermonde matrix for Scheme V.

	Scheme T	Scheme P	Scheme V
Communication Cost between User and Storage Node per Edit	$(\ell-1)\log q$ (worst case)	$\log \ell + \log q + 1$	$\log \ell + \log q$
Storage Overhead per Edit	0	$\log \ell$	$\log \ell$
Applicability	—	—	$\ell \leq q$ and for deletions only
Total Size of Storage Nodes for Data Blocks of Length ℓ'	$n\alpha\ell'$	$n\alpha\ell$, where ℓ is the length of the original data blocks	$n\alpha\ell'$

Table 2: Trade-off between storage overhead and communication complexity of various synchronization protocols.

the length of the original data file, while the latter scales storage requirements according to the number of unedited symbols.

Hybrid Schemes. Suppose first that we are given a constraint Γ on the storage overhead per edit, say $0 < \Gamma < \log \ell$. By combining Scheme V and Scheme T, we demonstrate a scheme that achieves a lower (worst case) communication cost while satisfying the aforementioned storage constraint.

Pick γ so that $\log(\gamma\ell) \leq \Gamma$. Next, divide each data block into two parts of lengths approximately $\gamma\ell$ and $(1-\gamma)\ell$. If the edit belongs to the first part, proceed with Scheme V by sending out $\log \ell + \log q$ bits and storing $\log(\gamma\ell)$ bits. On the other hand, if the edit belongs to the second part, then simply send the entire second part of $(1-\gamma)\ell \log q$ bits and the deleted position of $\log \ell$ bits with no storage overhead. Hence, in the worst case, the hybrid scheme incurs a storage overhead of $\max\{\log(\gamma\ell), 0\} = \log(\gamma\ell)$ bits and a communication cost of $\max\{\log(\gamma\ell) + \log q, (1-\gamma)\ell \log q\}$ bits. The choice of γ that minimizes the communication cost is given by the next lemma.

Lemma 6. *The hybrid scheme has smallest communication cost if*

$$\gamma = \min\left\{\frac{W(q^{\ell-1}\log q)}{\ell \log q}, \frac{2^{\Gamma}}{\ell}\right\}, \tag{12}$$

where $W(x)$ denotes the Lambert function, defined via $x = W(x)\exp(W(x))$.

Proof. The result easily follows by observing that $\max\{\log(\gamma\ell) + \log q, (1-\gamma)\ell \log q\}$ is minimized when $\log(\gamma\ell) + \log q = (1-\gamma)\ell \log q$. By rearranging terms, we arrive at the expression $z = y \exp y$ with $y = \gamma\,\ell \log q$ and $z = q^{\ell+1} \log q$, which yields the Lambert function form. The proof follows by observing that γ is also required to satisfy $\log(\gamma\,\ell) \leq \Gamma$. □

The hybrid scheme is described in detail in Appendix A and Figure 2 illustrates the communication cost and storage overhead trade-offs achievable by the hybrid scheme. Observe that both Scheme P and Scheme V come close to the lower bound when one is allowed a storage overhead of $\log \ell$ bits, assuming single edits per data block. Note that the lower bound on the communication cost is independent on the storage overhead. It would hence be of interest to derive a lower bound that actually captures the dependency on the storage overhead.

We conclude this section by considering a cost that takes into account *both* communication complexity and storage overhead. Similar to the hybrid scheme, we request certain users to communicate via Scheme P/V and others via Scheme T. Our new objective is to find an assignment

Figure 2: Communication Cost and Storage Overhead Tradeoff

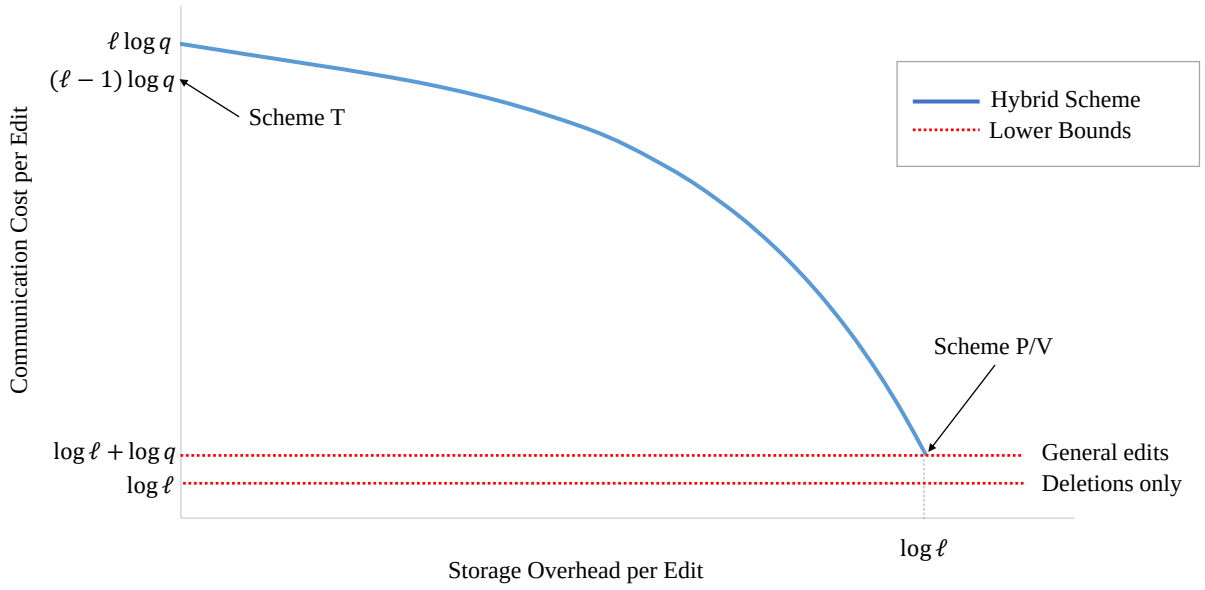

that minimizes the aggregate cost of storage and communication, rather than an assignment that obeys a storage overhead bound only.

Specifically, we consider a probabilistic edit model, where one edit occurs in each round and the probability of data block s being edited equals p_s. Let V be the set of users that are assigned Scheme P/V and define $p \triangleq \sum_{s \in V} p_s$.

Then, the expected storage overhead is $p \log \ell$, while the expected communication cost is $\log \ell + p \log q + (1-p)\ell \log q$. For some predefined $\theta \geq 0$, the aggregate cost is defined as

$$C_A(V) \triangleq \log \ell + p \log q + (1-p)\ell \log q + \theta p \log \ell. \tag{13}$$

Proposition 7. *Let C_A be given by (13). If $\theta \leq (\ell-1)\log q / \log \ell$, then C_A is minimized when $V = [B]$. Otherwise, if $\theta > (\ell-1)\log q / \log \ell$, then C_A is minimized when $V = \varnothing$.*

Proof. Consider $\theta \leq (\ell-1)\log q / \log \ell$. Suppose $V \neq [B]$ and pick $s \notin V$. Let $V' = V \cup \{s\}$. Then

$$C_A(V') - C_A(V) = p_s \log q - p_s \ell \log q + \theta p_s \log \ell = p_s(\theta \log \ell - (\ell-1)\log q) \leq 0.$$

Hence, augmenting V with s lowers the aggregate cost and so C_A is minimized when $V = [B]$. The argument for the other case proceeds along similar lines. □

6 Unknown Deletion Locations

In the previous sections, we considered synchronization protocols for user induced edits. In this case, the positions of deletions and the values of the corresponding symbols are known and available

to the users of the DSS. Many applications, including Dropbox-like file synchronization, call for file updates when changes are made by secondary users, in which case the edits are unknown to primary or other file users. As will be shown next, it is straightforward to accommodate the update scenarios to this model, provided that an additional small storage overhead is allowed. Again, we focus on the single edit scenario, for which we show that one only needs to store the Varshamov-Tenengolts (VT) Syndrome [45] of the data, in addition to a properly encoded original file. The single edit scenario applies to the case when frequent checks or updates are performed. When the updates appear less frequently, one may apply multiple-deletion correcting codes akin to those described in [46].

Consider the data string $\mathbf{x}^{(s)}$. Recall that its VT syndrome is given by

$$\nu_1^{(s)} \triangleq \sum_{i=1}^{\ell} x_i^{(s)},$$
$$\nu_2^{(s)} \triangleq \sum_{i=1}^{\ell-1} i\chi_i \bmod \ell,$$

where χ_i is the indicator function for the event $x_i^{(s)} \leq x_{i+1}^{(s)}$, $i \in [\ell - 1]$.

Case when all users have a single deletion. In addition to the matrices $\mathbf{A}^{(s)}$, we have each user store its VT syndrome $(\nu_1^{(s)}, \nu_2^{(s)})$, which is of size $\log q + \log \ell$ bits. Then prior to an update, each user retrieves its VT syndrome to compute the deleted position and its value. With this information, the user proceeds with any of the previously outlined update schemes.

Case where a proportion of data blocks have a deletion. To simplify arguments, we consider a systematic $[n, k]$ MDS code and the argument for a systematic $[n, k, d]$ regenerating code can be derived as a straightforward extension of the former case.

Observe that in the previous scheme, an additional overhead of $k(\log q + \log \ell)$ bits was required to store all VT syndromes for all k users. We next show that by using the structure of the MDS code, we can achieve storage savings by recording the VT syndromes of the $(n - k)$ check nodes only. For this purpose, assume that at most $n - k$ data blocks have a deletion and that $n - k < k$. Also, for simplicity, let $\ell = q$.

For the syndromes $\nu_1^{(1)}, \nu_1^{(2)}, \ldots, \nu_1^{(k)}$, each check node stores an additional check value that represents a linear combination of these syndromes, so that from any collection of k nodes all syndromes can be recovered. Similarly, each check node stores another check value from the syndromes[7] $\nu_2^{(1)}, \nu_2^{(2)}, \ldots, \nu_2^{(k)}$. Therefore, in this scheme, we store an additional $2(n-k)\log q < k(\log q + \log \ell)$ bits. Suppose that $n - k$ affected data blocks have a deletion. Each affected user requests the VT syndromes from the unaffected $k - (n - k)$ users and also the coded VT syndromes from the $n - k$ check nodes. With the k values, the affected user computes its own VT syndrome and consequently uses it to determine the position of its deletion and its value.

7 Conclusion

We presented a collection of protocols for synchronization of data encoded using regenerating and repair distributed storage codes, under probabilistic and deterministic symbol deletion edits.

[7]Here, we assume q is a prime and the syndrome is computed modulo q.

Our protocols trade communication overhead for small increases in storage overheads, and may be applied in various data deduplication and system update scenarios. The gist of the approach introduced was to use intermediary encoding, which allows the updated encoded information to retain repair and reconstruction properties, although with a different encoding functionality. In this case, the synchronization protocols may be seen as a form of functional update rules – rules that preserve functional properties but not the code structure. For the presented protocols, we provided simple fundamental lower bounds, average case performance bounds as well as an average case complexity analysis.

References

[1] F. Chang, J. Dean, S. Ghemawat, W. C. Hsieh, D. A. Wallach, M. Burrows, T. Chandra, A. Fikes, and R. E. Gruber, "Bigtable: A distributed storage system for structured data," *ACM Transactions on Computer Systems (TOCS)*, vol. 26, no. 2, p. 4, 2008.

[2] K. Shvachko, H. Kuang, S. Radia, and R. Chansler, "The hadoop distributed file system," in *Mass Storage Systems and Technologies (MSST), 2010 IEEE 26th Symposium on*, pp. 1–10, IEEE, 2010.

[3] B. Calder, J. Wang, A. Ogus, N. Nilakantan, A. Skjolsvold, S. McKelvie, Y. Xu, S. Srivastav, J. Wu, H. Simitci, *et al.*, "Windows azure storage: a highly available cloud storage service with strong consistency," in *Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles*, pp. 143–157, ACM, 2011.

[4] W. Zeng, Y. Zhao, K. Ou, and W. Song, "Research on cloud storage architecture and key technologies," in *Proceedings of the 2nd International Conference on Interaction Sciences: Information Technology, Culture and Human*, pp. 1044–1048, ACM, 2009.

[5] S. B. Wicker and V. K. Bhargava, *Reed-Solomon codes and their applications*. John Wiley & Sons, 1999.

[6] A. G. Dimakis, P. B. Godfrey, Y. Wu, M. J. Wainwright, and K. Ramchandran, "Network coding for distributed storage systems," *Information Theory, IEEE Transactions on*, vol. 56, no. 9, pp. 4539–4551, 2010.

[7] C. Huang, H. Simitci, Y. Xu, A. Ogus, B. Calder, P. Gopalan, J. Li, S. Yekhanin, *et al.*, "Erasure coding in windows azure storage.," in *USENIX Annual Technical Conference*, pp. 15–26, 2012.

[8] K. Rashmi, N. B. Shah, P. V. Kumar, and K. Ramchandran, "Explicit construction of optimal exact regenerating codes for distributed storage," in *Communication, Control, and Computing, 2009. Allerton 2009. 47th Annual Allerton Conference on*, pp. 1243–1249, IEEE, 2009.

[9] P. Gopalan, C. Huang, H. Simitci, and S. Yekhanin, "On the locality of codeword symbols," *Information Theory, IEEE Transactions on*, vol. 58, no. 11, pp. 6925–6934, 2012.

[10] P. Gopalan, C. Huang, B. Jenkins, and S. Yekhanin, "Explicit maximally recoverable codes with locality.," in *Electronic Colloquium on Computational Complexity (ECCC)*, vol. 20, p. 104, 2013.

[11] K. A. S. Immink, *Codes for mass data storage systems.* Shannon Foundation Publisher, 2004.

[12] S. Gregori, A. Cabrini, O. Khouri, and G. Torelli, "On-chip error correcting techniques for new-generation flash memories," *Proceedings of the IEEE*, vol. 91, no. 4, pp. 602–616, 2003.

[13] E. Yaakobi, J. Ma, L. Grupp, P. H. Siegel, S. Swanson, and J. K. Wolf, "Error characterization and coding schemes for flash memories," in *GLOBECOM Workshops (GC Wkshps), 2010 IEEE*, pp. 1856–1860, IEEE, 2010.

[14] S. A. Weil, S. A. Brandt, E. L. Miller, D. D. Long, and C. Maltzahn, "Ceph: A scalable, high-performance distributed file system," in *Proceedings of the 7th symposium on Operating systems design and implementation*, pp. 307–320, USENIX Association, 2006.

[15] D. Meister and A. Brinkmann, "Multi-level comparison of data deduplication in a backup scenario," in *Proceedings of SYSTOR 2009: The Israeli Experimental Systems Conference*, p. 8, ACM, 2009.

[16] I. Drago, M. Mellia, M. M Munafo, A. Sperotto, R. Sadre, and A. Pras, "Inside dropbox: understanding personal cloud storage services," in *Proceedings of the 2012 ACM conference on Internet measurement conference*, pp. 481–494, ACM, 2012.

[17] P. Riteau, C. Morin, and T. Priol, "Shrinker: Improving live migration of virtual clusters over wans with distributed data deduplication and content-based addressing," in *Euro-Par 2011 Parallel Processing*, pp. 431–442, Springer, 2011.

[18] Y. Zhang, Y. Wu, and G. Yang, "Droplet: a distributed solution of data deduplication," in *Proceedings of the 2012 ACM/IEEE 13th International Conference on Grid Computing*, pp. 114–121, IEEE Computer Society, 2012.

[19] A. Tridgell, P. Mackerras, *et al.*, "The rsync algorithm," 1996.

[20] T. Knauth and C. Fetzer, "dsync: Efficient block-wise synchronization of multi-gigabyte binary data.," in *LISA*, pp. 45–58, 2013.

[21] A. Orlitsky and K. Viswanathan, "One-way communication and error-correcting codes," *Information Theory, IEEE Transactions on*, vol. 49, no. 7, pp. 1781–1788, 2003.

[22] Y. Minsky, A. Trachtenberg, and R. Zippel, "Set reconciliation with nearly optimal communication complexity," *Information Theory, IEEE Transactions on*, vol. 49, no. 9, pp. 2213–2218, 2003.

[23] R. Venkataramanan, H. Zhang, and K. Ramchandran, "Interactive low-complexity codes for synchronization from deletions and insertions," in *Communication, Control, and Computing (Allerton), 2010 48th Annual Allerton Conference on*, pp. 1412–1419, IEEE, 2010.

[24] N. Ma, K. Ramchandran, and D. Tse, "Efficient file synchronization: A distributed source coding approach," in *Information Theory Proceedings (ISIT), 2011 IEEE International Symposium on*, pp. 583–587, IEEE, 2011.

[25] S. S. T. Yazdi and L. Dolecek, "Synchronization from deletions through interactive communication," in *Turbo Codes and Iterative Information Processing (ISTC), 2012 7th International Symposium on*, pp. 66–70, IEEE, 2012.

[26] N. Bitouze and L. Dolecek, "Synchronization from insertions and deletions under a non-binary, non-uniform source," in *Information Theory Proceedings (ISIT), 2013 IEEE International Symposium on*, pp. 2930–2934, IEEE, 2013.

[27] R. Venkataramanan, V. N. Swamy, and K. Ramchandran, "Efficient interactive algorithms for file synchronization under general edits," in *Communication, Control, and Computing (Allerton), 2013 51st Annual Allerton Conference on*, 2013.

[28] S. Tabatabaei Yazdi and L. Dolecek, "A deterministic, polynomial-time protocol for synchronizing from deletions," *Information Theory, IEEE Transactions on*, vol. 60, no. 1, pp. 397–409, 2014.

[29] N. P. Anthapadmanabhan, E. Soljanin, and S. Vishwanath, "Update-efficient codes for erasure correction," in *Communication, Control, and Computing (Allerton), 2010 48th Annual Allerton Conference on*, pp. 376–382, IEEE, 2010.

[30] A. S. Rawat, S. Vishwanath, A. Bhowmick, and E. Soljanin, "Update efficient codes for distributed storage," in *Information Theory Proceedings (ISIT), 2011 IEEE International Symposium on*, pp. 1457–1461, IEEE, 2011.

[31] A. Mazumdar, G. W. Wornell, and V. Chandar, "Update efficient codes for error correction," in *Information Theory Proceedings (ISIT), 2012 IEEE International Symposium on*, pp. 1558–1562, IEEE, 2012.

[32] A. Mazumdar, V. Chandar, and G. W. Wornell, "Update-efficiency and local repairability limits for capacity approaching codes," *Selected Areas in Communications, IEEE Journal on*, vol. 32, no. 5, pp. 976–988, 2014.

[33] F. J. MacWilliams and N. J. A. Sloane, *The theory of error-correcting codes*, vol. 16. Elsevier, 1977.

[34] Y. Wu, A. G. Dimakis, and K. Ramchandran, "Deterministic regenerating codes for distributed storage," in *Proc. Allerton Conference on Control, Computing and Communication*, 2007.

[35] K. Rashmi, N. B. Shah, P. V. Kumar, and K. Ramchandran, "Explicit construction of optimal exact regenerating codes for distributed storage," in *Allerton Conference on Control, Computing, and Communication*, 2009.

[36] A. S. Rawat, O. O. Koyluoglu, N. Silberstein, and S. Vishwanath, "Optimal locally repairable and secure codes for distributed storage systems," in *arXiv:1210.6954*, 2013.

[37] S. Arora and B. Barak, *Computational Complexity: A Modern Approach*. New York, NY, USA: Cambridge University Press, 1st ed., 2009.

[38] V. I. Levenshtein, "Efficient reconstruction of sequences from their subsequences or supersequences," *Journal of Combinatorial Theory, Series A*, vol. 93, no. 2, pp. 310–332, 2001.

[39] A. Orlitsky, *Communication issues in distributed computing*. PhD thesis, Stanford University, 1987.

[40] V. I. Levenshtein, "Binary codes capable of correcting deletions, insertions and reversals," in *Soviet physics doklady*, vol. 10, p. 707, 1966.

[41] L. Calabi and W. Hartnett, "Some general results of coding theory with applications to the study of codes for the correction of synchronization errors," *Information and Control*, vol. 15, no. 3, pp. 235–249, 1969.

[42] D. S. Hirschberg and M. Regnier, "Tight bounds on the number of string subsequences," *Journal of Discrete Algorithms*, vol. 1, no. 1, pp. 123–132, 2000.

[43] H. Mercier, M. Khabbazian, and V. K. Bhargava, "On the number of subsequences when deleting symbols from a string," *IEEE transactions on information theory*, vol. 54, no. 7, pp. 3279–3285, 2008.

[44] Y. Liron and M. Langberg, "A characterization of the number of subsequences obtained via the deletion channel," in *Information Theory Proceedings (ISIT), 2012 IEEE International Symposium on*, pp. 503–507, IEEE, 2012.

[45] N. J. Sloane, "On single-deletion-correcting codes," *Codes and Designs, de Gruyter, Berlin*, pp. 273–291, 2002.

[46] M. C. Davey and D. J. MacKay, "Reliable communication over channels with insertions, deletions, and substitutions," *Information Theory, IEEE Transactions on*, vol. 47, no. 2, pp. 687–698, 2001.

[47] H. A. David and H. N. Nagaraja, *Order statistics*. Wiley Online Library, 1970.

[48] N. Balakrishnan, C. Charalambides, and N. Papadatos, "Bounds on expectation of order statistics from a finite population," *Journal of statistical planning and inference*, vol. 113, no. 2, pp. 569–588, 2003.

A Hybrid Schemes

We describe in detail the hybrid scheme that combines Scheme P and Scheme V. Define $\ell^* = (1-\gamma)\ell$. Then, for $s \in [B]$, initialize

$$\mathbf{A}^{(s)} \leftarrow \begin{pmatrix} \mathbf{V} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{pmatrix},$$

where $\mathbf{V}$ is a Vandermonde matrix of dimension $(\ell - \ell^*) \times (\ell - \ell^*)$ and $\mathbf{I}$ is an identity matrix of dimension ℓ^*. For subsequent edits, we proceed as below.

Scheme H: Assumption: $\mathbf{x}^{(s)}$ is deleted at coordinate i_s.

if $i_s \leq \ell - \ell^*$ **then**
  User s sends to all connected storage nodes $N(s)$ the value of the deleted symbol – $x_{i_s}^{(s)}$ ($\log q$ bits) and the coordinate i_s ($\log \ell$ bits)
  $\mathbf{A}^{(s)} \leftarrow \mathbf{A}^{(s)}$ with the i_sth row removed
else
  User s computes the vector
  $\mathbf{d} \leftarrow \left(x_{\ell-\ell^*+1}^{(s)}, x_{\ell-\ell^*+2}^{(s)}, \ldots, x_{\ell}^{(s)}\right) - \left(x_{\ell-\ell^*+1}^{(s)}, \ldots, x_{i_s-1}^{(s)}, x_{i_s+1}^{(s)}, \ldots, x_{\ell}^{(s)}, 0\right)$
  User s sends to all connected storage nodes $N(s)$ the vector $\mathbf{d}$ ($\ell^* \log q$ bits)
  User s pads $\mathbf{x}^{(s)}$ with a zero.
end
for $t \in N(s)$ **do**
  Subtract $\mathbf{D}|_{\{t\}\times[\alpha]\times[\ell]}$ where

$$\mathbf{D} = \begin{cases} \textsc{Encode}(\ell)\left(\mathbf{0}, \mathbf{0}, \ldots, x_{i_s}^{(s)}\mathbf{A}|_{\{i_s\}\times[\ell]}, \ldots, \mathbf{0}\right), & \text{if } i_s \leq \ell - \ell^*, \\ \textsc{Encode}(\ell)\left(\mathbf{0}, \mathbf{0}, \ldots, (\mathbf{0}, \mathbf{d}), \ldots, \mathbf{0}\right), & \text{otherwise.} \end{cases}$$

end

Proposition 8 (Scheme H)**.** *Consider an $(n, k, d, \alpha\ell, B\ell)$ DSS code and assume a single deletion per user. The updates in accordance to Scheme H result in an $(n, k, d, \alpha\ell, B\ell)$ DSS code. In the worst case, each user introduces a storage overhead of $\max\{\log(\gamma\ell), 0\}$ bits and incurs a communication cost of $\max\{\log(\gamma\ell) + \log q, (1-\gamma)\ell \log q\}$ bits for its edit.*

The proof is similar to the proofs of Proposition 4 and Proposition 5 and therefore omitted.

Example 5. We illustrate Scheme H via a simple example involving a $[3, 2]$ MDS code over $\mathbb{F}_5$, involving two arbitrarily chosen data blocks $\mathbf{u}^{(1)}$ and $\mathbf{u}^{(2)}$ shown below.

$\mathbf{u}^{(1)}$	$\mathbf{u}^{(2)}$	$\mathbf{D}$	$\textsc{Encode}(\ell)\left(\mathbf{u}^{(1)}\mathbf{A}^{(1)}, \mathbf{u}^{(2)}\mathbf{A}^{(2)}\right)$	$\mathbf{A}^{(1)}$	$\mathbf{A}^{(2)}$
$(1,1,1,1,1,1,1)$	$(1,2,3,4,3,2,1)$	—	$\begin{pmatrix} 4&0&0&0&1&1&1 \\ 0&0&0&4&3&2&1 \\ 4&0&0&4&4&3&2 \end{pmatrix}$	$\begin{pmatrix} 1&1&1&1&0&0&0 \\ 1&2&4&3&0&0&0 \\ 1&3&4&2&0&0&0 \\ 1&4&1&4&0&0&0 \\ 0&0&0&0&1&0&0 \\ 0&0&0&0&0&1&0 \\ 0&0&0&0&0&0&1 \end{pmatrix}$	$\begin{pmatrix} 1&1&1&1&0&0&0 \\ 1&2&4&3&0&0&0 \\ 1&3&4&2&0&0&0 \\ 1&4&1&4&0&0&0 \\ 0&0&0&0&1&0&0 \\ 0&0&0&0&0&1&0 \\ 0&0&0&0&0&0&1 \end{pmatrix}$
$(1,1,\not{1},1,1,1,1)$	$(1,2,3,4,3,2,1)$	$\begin{pmatrix} 1&3&4&2&0&0&0 \\ 0&0&0&0&0&0&0 \\ 1&3&4&2&0&0&0 \end{pmatrix}$	$\begin{pmatrix} 3&2&1&3&1&1&1 \\ 0&0&0&4&3&2&1 \\ 3&2&1&2&4&3&2 \end{pmatrix}$	$\begin{pmatrix} 1&1&1&1&0&0&0 \\ 1&2&4&3&0&0&0 \\ 1&4&1&4&0&0&0 \\ 0&0&0&0&1&0&0 \\ 0&0&0&0&0&1&0 \\ 0&0&0&0&0&0&1 \end{pmatrix}$	$\begin{pmatrix} 1&1&1&1&0&0&0 \\ 1&2&4&3&0&0&0 \\ 1&3&4&2&0&0&0 \\ 1&4&1&4&0&0&0 \\ 0&0&0&0&1&0&0 \\ 0&0&0&0&0&1&0 \\ 0&0&0&0&0&0&1 \end{pmatrix}$
$(1,1,1,1,1,1)$	$(1,2,3,4,\not{3},2,1)$	$\begin{pmatrix} 0&0&0&0&0&0&0 \\ 0&0&0&0&1&1&1 \\ 0&0&0&0&1&1&1 \end{pmatrix}$	$\begin{pmatrix} 3&2&1&3&1&1&1 \\ 0&0&0&4&2&1&0 \\ 3&2&1&2&3&2&1 \end{pmatrix}$	$\begin{pmatrix} 1&1&1&1&0&0&0 \\ 1&2&4&3&0&0&0 \\ 1&4&1&4&0&0&0 \\ 0&0&0&0&1&0&0 \\ 0&0&0&0&0&1&0 \\ 0&0&0&0&0&0&1 \end{pmatrix}$	$\begin{pmatrix} 1&1&1&1&0&0&0 \\ 1&2&4&3&0&0&0 \\ 1&3&4&2&0&0&0 \\ 1&4&1&4&0&0&0 \\ 0&0&0&0&1&0&0 \\ 0&0&0&0&0&1&0 \\ 0&0&0&0&0&0&1 \end{pmatrix}$
$(1,1,1,1,1,1)$	$(\not{1},2,3,4,2,1,0)$	$\begin{pmatrix} 0&0&0&0&0&0&0 \\ 1&1&1&1&0&0&0 \\ 1&1&1&1&0&0&0 \end{pmatrix}$	$\begin{pmatrix} 3&2&1&3&1&1&1 \\ 4&4&4&3&2&1&0 \\ 2&1&0&1&3&2&1 \end{pmatrix}$	$\begin{pmatrix} 1&1&1&1&0&0&0 \\ 1&2&4&3&0&0&0 \\ 1&4&1&4&0&0&0 \\ 0&0&0&0&1&0&0 \\ 0&0&0&0&0&1&0 \\ 0&0&0&0&0&0&1 \end{pmatrix}$	$\begin{pmatrix} 1&2&4&3&0&0&0 \\ 1&3&4&2&0&0&0 \\ 1&4&1&4&0&0&0 \\ 0&0&0&0&1&0&0 \\ 0&0&0&0&0&1&0 \\ 0&0&0&0&0&0&1 \end{pmatrix}$

We reconstruct the data blocks using nodes or rows 1 and 3.

1. From row 1 and 3, we infer that row 2 is $(4, 4, 4, 3, 2, 1, 0)$.
2. Hence, we obtain

$$\hat{\mathbf{u}}^{(1)} = (3,2,1,3,1,1,1)\left(\mathbf{A}^{(1)}\right)^{-1} = (1,1,1,1,1,1),$$
$$\hat{\mathbf{u}}^{(2)} = (4,4,4,3,2,1,0)\left(\mathbf{A}^{(2)}\right)^{-1} = (2,3,4,2,1,0),$$

as desired.

B Probabilistic Analysis of the Communication Cost

In Section 5.2, we described a straightforward worst case analysis of the communication costs incurred by various synchronization/update protocols. In what follows, we estimate the expected communication costs of the schemes in the asymptotic regime, and in particular, show that even on average, Scheme T requires communicating $\Omega(\ell \log q)$ bits per one single round of editing, unlike schemes P and V.

For the purpose of average case analysis, we assume the following probabilistic models for edits:

(UD) The uniform deletions model. For simplicity, the underlying assumption is that each data block has a single deletion, at a coordinate chosen uniformly at random.

(CND) The combinatorial nonuniform deletion model. Here, the model assumes that exactly D deletions occurred uniformly at random over all the ℓ coordinates of the B users. In addition, we assume $D = o(B\ell/\log \ell)$.

(PND) The probabilistic nonuniform deletion model. Here, we assume that each of the ℓ coordinates of the B users is *independently* edited with probability p, where p may depend on ℓ. In addition, we assume that the edit probability is relatively small, i.e., that $p = o(1/\log \ell)$. Hence, the total number of edits in a data block is $o(\ell/\log \ell)$.

Theorem 9 (Scheme T)**.** *Consider Scheme T. Let C be the communication cost between a user and a storage node, with edits occurring in accordance with the (UD), (CND) and (PND) models. Then*

$$E[C] = \eta(\ell)\, \ell \log q,$$

where

1. *For model (UD), we have*

$$\lim_{\ell\to\infty} \eta(\ell) \begin{cases} = \dfrac{B-1}{B+1}, & \text{if } B \text{ is constant,} \quad (14a) \\ \geq 1 - 2e^{-c}, & \text{if } \lim_{\ell\to\infty} B(\ell)/\ell = c. \quad (14b) \end{cases}$$

2. *For model (CND), we have*

$$\lim_{\ell\to\infty} \eta(\ell) \geq \frac{D}{D+1}. \quad (15)$$

3. *For model (PND), we have*

$$\lim_{\ell\to\infty} \eta(\ell) \geq 1 - e^{-c}, \quad \text{if } \lim_{\ell\to\infty} B(\ell)\, p(\ell) = c. \quad (16)$$

Proof. Recall the definitions of $i_{\max}$, $i_{\min}$ and I from Section 3. For the model (UD), we have $|I| = i_{\max} - i_{\min}$, while $|I| = \ell - i_{\min}$ holds for models (CND) and (PND). Hence, $\eta(\ell)\ell = E[|I|] = E[i_{\max}] - E[i_{\min}] = \ell - 2E[i_{\min}]$ for model (UD), and $\eta(\ell)\ell = \ell - E[i_{\min}]$ for models (CND) and (PND). Therefore, our problem reduces to estimating $E[i_{\min}]$ for the various models. To do so, we follow the standard derivation methods in order statistics (for example, see David and Nagaraja [47]), and point out that the range (span, support) of the observed deletion positions is studied in order statistics under the name *sample range*.

Recall that $E[i_{\min}] = \sum_{i=1}^{\ell} \text{Prob}(i_{\min} \geq i)$ and observe that

$$\text{Prob}(i_{\min} \geq i) = \begin{cases} \left(\frac{\ell-i+1}{\ell}\right)^B, & \text{for model (UD),} \\ \frac{\binom{B(\ell-i+1)}{D}}{\binom{B\ell}{D}}, & \text{for model (CND)} \\ (1-p)^{B(i-1)}, & \text{for model (PND).} \end{cases}$$

Hence, for model (UD), we have $E[i_{\min}]/\ell = \sum_{i=1}^{\ell} i^B/\ell^{B+1}$. When B is constant, $E[i_{\min}]/\ell \to 1/(B+1)$ as $\ell \to \infty$ yielding (14a). When $\lim_{\ell\to\infty} B/\ell = \kappa$, we have

$$\frac{E[i_{\min}]}{\ell} \leq \frac{1}{\ell} + (\ell-1)\frac{(\ell-1)^B}{\ell^{B+1}} = \frac{1}{\ell} + \left(1 - \frac{1}{\ell}\right)^{B+1} \to e^{-\kappa},$$

yielding (14b).

On the other hand, for model (CND), we have[8]

$$E[i_{\min}] = \frac{\sum_{i=1}^{\ell} \binom{Bi}{D}}{\binom{B\ell}{D}} \leq \frac{\sum_{i=D}^{B\ell} \binom{i}{D} + (B-1)\binom{B\ell}{D}}{B\binom{B\ell}{D}} = \frac{\binom{B\ell+1}{D+1}}{B\binom{B\ell}{D}} + \frac{B-1}{B} = \frac{B\ell+1}{B(D+1)} + \frac{B-1}{B}$$

As a result, $\lim_{\ell\to\infty} E[i_{\min}]/\ell \leq 1/(D+1)$, yielding (15).

Finally, for model (PND), we have

$$E[i_{\min}] = \sum_{i=1}^{\ell} (1-p)^{B(i-1)\ell} \leq 1 + (\ell-1)(1-p)^B.$$

If $\lim_{\ell\to\infty} Bp = c$, then $\lim_{\ell\to\infty} E[i_{\min}]/\ell \leq e^{-c}$ yields (16).

□

Proposition 10 (Scheme P/V)**.** *Consider either Scheme P or Scheme V. Define C to be the communication cost between one selected user and a connected storage node, with edits occurring according to the above models. Then*

1. *For model (UD),*
$$E[C] = \log \ell + \log q.$$
2. *For model (CND),*
$$E[C] = \frac{D}{B}(\log \ell + \log q).$$
3. *For model (PND),*
$$E[C] = p\ell(\log \ell + \log q).$$

Proof. Note that the communication cost between a user and a connected storage node for Schemes P/V is given by $d(\log q + \log \ell)$, where d denotes the corresponding number of edits of the user. Since for the models (UD), (CND) and (PND), the expected number of edits are one, D/B and $p\ell$, respectively, the proposition follows after straightforward algebraic manipulations. □

In Appendix C, we explored other methods for reducing the expected communication cost of the Schemes P/V. The idea behind the method is to observe the location of the edit with largest index, and then communicate all deletion positions based on a rescaled data block length equal to that index. As an illustration, if the largest index of an edit in a data block of length ℓ is of the order $o(\ell)$, then the locations of the deletions would only require $o(\log n)$ bits for encoding. For the uniform and at random deletion models, these savings are unfortunately not significant.

We conclude this analysis by comparing the relationships between the expected communication costs of the various models. Let C_T and C_P be the expected communication cost in Scheme T and Scheme P/V, respectively. In the following analysis, asymptotics are computed in ℓ and we assume that q is either fixed or that $q = O(\ell)$.

We note that under our model assumptions, we have $\lim_{\ell\to\infty} C_P/C_T = 0$.

[8] A similar formulation of $E[i_{\min}]$ can be found in [48].

1. For model (UD), we have $C_T = \Omega(\ell \log q)$, while $C_P = \Theta(\log \ell + \log q)$. Hence, $\lim_{\ell\to\infty} C_P/C_T = 0$.

2. For model (CND), we have $C_T = \Omega(\ell \log q)$ or $C_T \geq \kappa \ell \log q$, for some constant κ. Then

$$\lim_{\ell\to\infty} \frac{C_P}{C_T} \leq \lim_{\ell\to\infty} \frac{D/B(\log \ell + \log q)}{\kappa \ell \log q} \to 0,$$

since $D = o(B\ell/\log \ell)$.

3. For model (PND), we have $C_T = \Omega(\ell \log q)$ or $C_T \geq \kappa \ell \log q$, for some constant κ. Then

$$\lim_{\ell\to\infty} \frac{C_P}{C_T} \leq \lim_{\ell\to\infty} \frac{p\ell(\log \ell + \log q)}{\kappa \ell \log q} \to 0,$$

since $p = o(1/\log \ell)$.

Consequently, for all three probabilistic models and our proposed protocols, we achieve a communication complexity that is asymptotically negligible compared to that of the traditional approach given by Scheme T.

C Reducing the Expected Communication Cost for Schemes P/V

Suppose that the largest edited coordinate in a user data block with d edits equals $i_{\max}$. In this case, edits are confined to a shorter block than the length of the data block, and sending $d\lceil \log i_{\max} \rceil$ bits of information for d edits suffices. In what follows, we compute the expected value of $\lceil \log i_{\max} \rceil$ given d edits, in order to determine how much one may save in the communication rate due to this simple observation.

First, observe that for the independent, uniform deletion model, we have

$$\begin{aligned} E[\lceil \log i_{\max} \rceil \text{ given } d \text{ edits}] &= \sum_{i=1}^{\lceil \log \ell \rceil} \text{Prob}(\lceil \log i_{\max} \rceil \geq i) = \sum_{i=1}^{\lceil \log \ell \rceil} \text{Prob}(i_{\max} \geq 2^{i-1} + 1) \\ &= \sum_{i=1}^{\lceil \log \ell \rceil} \left(1 - \frac{\binom{2^{i-1}}{d}}{\binom{\ell}{d}} \right) = \lceil \log \ell \rceil - \sum_{i=1}^{\lceil \log \ell \rceil} \frac{\binom{2^{i-1}}{d}}{\binom{\ell}{d}}. \end{aligned}$$

Let $s(\ell, d) \triangleq \sum_{i=1}^{\lceil \log \ell \rceil} \binom{2^{i-1}}{d} / \binom{\ell}{d}$ be the second term in the previously derived expression, representing the expected number of bits that can be saved by using the reduced range encoding scheme. It is straightforward to see that $\sum_{i=1}^{\lceil \log \ell \rceil} \binom{2^{i-1}}{d} / \binom{\ell}{d} < \sum_{i=1}^{\lceil \log \ell \rceil} 2^{i-1}/\ell \leq (2\ell - 1)/\ell < 2$.

Assuming model (CND), the expected communication cost per user-storage node pair follows from a generating function approach that gives

$$\begin{aligned} &\frac{1}{\binom{B+D-1}{D}} [\lambda^D](1-\lambda)^{-(B-1)} \sum_{d\geq 1} d\lambda^d (\lceil \log \ell \rceil - s(\ell, d)) \\ &= \frac{\lceil \log \ell \rceil}{\binom{B+D-1}{D}} [\lambda^D](1-\lambda)^{-(B-1)} \sum_{d\geq 1} d\lambda^d - \frac{1}{\binom{B+D-1}{D}} [\lambda^D](1-\lambda)^{-(B-1)} \sum_{d\geq 1} d\lambda^d s(\ell, d) \\ &= \frac{\lceil \log \ell \rceil}{\binom{B+D-1}{D}} [\lambda^D] \frac{\lambda}{(1-\lambda)^{B+1}} - \frac{1}{\binom{B+D-1}{D}} \sum_{d=1}^{D} d s(\ell, d) [\lambda^{D-d}](1-\lambda)^{-(B-1)} \end{aligned}$$

$$= \frac{\binom{B+D-1}{D-1}}{\binom{B+D-1}{D}} \lceil \log \ell \rceil - \frac{1}{\binom{B+D-1}{D}} \sum_{d=1}^{D} ds(\ell, d) \binom{B+D-d-2}{D-d}$$
$$= \frac{D}{B} \lceil \log \ell \rceil - \frac{1}{\binom{B+D-1}{D}} \sum_{d=1}^{D} \sum_{i=1}^{\lceil \log \ell \rceil} \frac{d\binom{2^{i-1}}{d}\binom{B+D-d-2}{D-d}}{\binom{\ell}{d}},$$

where $[x^k]p(x)$ denotes the coefficient of x^k in the polynomial $p(x)$.

Assuming model (PND), the expected communication cost may be evaluated as

$$\sum_{d=1}^{\ell} dp^d(1-p)^{\ell-d}\binom{\ell}{d}(\lceil \log \ell \rceil - s(\ell, d)) = \lceil \log \ell \rceil \sum_{d=1}^{\ell} dp^d(1-p)^{\ell-d}\binom{\ell}{d} - \sum_{d=1}^{\ell} dp^d(1-p)^{\ell-d}\binom{\ell}{d} s(\ell, d)$$
$$= p\ell \lceil \log \ell \rceil - \sum_{d=1}^{\ell} dp^d(1-p)^{\ell-d}\binom{\ell}{d} s(\ell, d) = p\ell \lceil \log \ell \rceil - \sum_{d=1}^{\ell} \sum_{i=1}^{\lceil \log \ell \rceil} dp^d(1-p)^{\ell-d}\binom{2^{i-1}}{d}.$$

Assuming model (CND) with different values of λ_s, the expected communication cost for user s is given by

$$\frac{\sum_{d_1+d_2+\cdots+d_B=D} \lambda_1^{d_1}\lambda_2^{d_2}\cdots\lambda_B^{d_B} d_s(\lceil \log \ell \rceil - s(\ell, d_s))}{\sum_{d_1+d_2+\cdots+d_B=D} \lambda_1^{d_1}\lambda^{d_2}\cdots\lambda_B^{d_B}}$$
$$= \frac{\sum_{d_s=0}^{D} \lambda_s^{d_s} d_s(\lceil \log \ell \rceil - s(\ell, d_s)) \sum_{\sum_{i\neq s} d_i = D-d_s} \lambda_1^{d_1}\lambda_2^{d_2}\cdots\lambda_{s-1}^{d_{s-1}}\lambda_{s+1}^{d_{s+1}}\cdots\lambda_B^{d_B}}{\sum_{d_1+d_2+\cdots+d_B=D} \lambda_1^{d_1}\lambda^{d_2}\cdots\lambda_B^{d_B}}$$
$$= \frac{\sum_{d_s=0}^{D} \lambda_s^{d_s} d_s(\lceil \log \ell \rceil - s(\ell, d_s)) h_{D-d_s}(\lambda_1, \lambda_2, \ldots, \lambda_{s-1}, \lambda_{s+1}, \ldots, \lambda_B)}{h_D(\lambda_1, \lambda_2, \ldots, \lambda_B)},$$

where $h_k(X_1, X_2, \ldots, X_n)$ is the *complete homogeneous symmetric polynomial of degree k in n variables.*

If $\lambda_s = \lambda^s$, we can specialize to

$$h_k(\lambda, \lambda^2, \ldots, \lambda^B) = [z^k] \prod_{1\leq i\leq B} (1-\lambda^i z)^{-1} \tag{17}$$
$$h_k(\lambda, \lambda^2, \ldots, \lambda^{s-1}, \lambda^{s+1}, \ldots, \lambda^B) = [z^k] \prod_{1\leq i\leq B, i\neq s} (1-\lambda^i z)^{-1}. \tag{18}$$

We remark that (17) yields the generating function that counts the number of partitions with at most k parts of size at most B, while (18) yields the generating function that counts the number of partitions with at most k parts of size at most B and no part of size s.